

A stochastic optimisation model to support cybersecurity within the UK national health service

Emilia Grass, Christina Pagel, Sonya Crowe & Saira Ghafur

To cite this article: Emilia Grass, Christina Pagel, Sonya Crowe & Saira Ghafur (17 Dec 2024): A stochastic optimisation model to support cybersecurity within the UK national health service, Journal of the Operational Research Society, DOI: [10.1080/01605682.2024.2436063](https://doi.org/10.1080/01605682.2024.2436063)

To link to this article: <https://doi.org/10.1080/01605682.2024.2436063>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



[View supplementary material](#)



Published online: 17 Dec 2024.



[Submit your article to this journal](#)



Article views: 17



[View related articles](#)



[View Crossmark data](#)

A stochastic optimisation model to support cybersecurity within the UK national health service

Emilia Grass^a , Christina Pagel^b , Sonya Crowe^b  and Saira Ghafur^c

^aKarlsruhe Institute of Technology, Karlsruhe, Germany; ^bUniversity College London, London, UK; ^cImperial College London, London, UK

ABSTRACT

Over the past decade, the adoption of new digital technologies in healthcare has surged, significantly enhancing care delivery and accessibility. However, this digital transformation has been accompanied by a sharp increase in cyber-attacks, posing severe risks to hospital functionality and patient safety. To address the challenge of planning for uncertain future cyber incidents, we propose a two-stage stochastic model designed to bolster the cyber resilience of healthcare providers by selecting optimal countermeasures in preparation for upcoming cyber incidents. Numerical tests demonstrate the model's effectiveness, with the Value of the Stochastic Solution showing a 21% improvement over a deterministic approach. To be optimally equipped even for low-probability high-impact attacks we incorporate the risk measure *Conditional Value-at-Risk*. The corresponding countermeasure solution led to 44% fewer rejected patients in a worst-case scenario. The robustness of the solution is underscored by its consistent performance across various scenarios, budget levels, and risk preferences, making it a reliable tool for enhancing cybersecurity in healthcare. These results highlight the importance of tailored, robust cybersecurity strategies in healthcare, ensuring preparedness for a wide range of potential threats.

ARTICLE HISTORY

Received 29 November 2023
Accepted 26 November 2024

KEYWORDS



Stochastic programming;
cybersecurity; risk measure
CVaR; health services; NHS


1. Introduction

In recent years, the number and severity of cyber-attacks against healthcare providers and hospitals has increased significantly. According to ITPro (2024) healthcare was one of the top three most targeted sectors in 2023, experiencing an average of 1500 weekly attacks. In May 2021, the Irish health system experienced its most widespread ransomware attack to date, where access to electronic systems and data was blocked, severely impacting critical services such as gynaecology and maternity clinics as well as cancer and children's care. When Change Healthcare, a major healthcare technology company in the United States, was hit by a cyber-attack in February 2024, over 100 hospitals across the country had their ability to bill for care hindered, impacting payroll and overall operations. Another recent cyber-attack took place in Romania in February 2024, where ransomware attacks led to the shutdown of 100 hospitals' digital systems. The attack started with a children's hospital and spread to other facilities, forcing hospitals to revert to manual paper records for patient admissions and medical recommendations. Additionally, the ransomware cyber-attack against pathology services provider

Synnovis in June 2024, targeted hospitals in London, resulting in the postponement of 9423 acute outpatient appointments and 1660 elective procedures. During the May 2017 WannaCry attack on the English NHS, hospitals were locked out of digital systems and medical devices like MRI scanners, severely limiting patient care. The shutdown of intranets and electronic records forced staff to use manual processes, cancel appointments, and divert emergency ambulances.

Employing appropriate countermeasures before an attack can significantly improve cybersecurity. However, deciding on the number and type of these countermeasures is difficult. Almost 1200 cybersecurity countermeasures are available (Federal Office for Information Security, 2023), encompassing anti-virus software, tools for cyber threat analysis and data loss prevention. This abundance complicates healthcare managers' task of selecting cost-effective and efficient tools for improving cybersecurity. Moreover, adhering to local regulations like the Cyber Essentials Plus (CEP) scheme and the General Data Protection Regulation (GDPR) adds complexity. Deciding on prevention strategies is further complicated by uncertainty—precise details

CONTACT Emilia Grass  emilia.grass@kit.edu  Karlsruhe Institute of Technology, Institute of Information Security and Dependability, Kaiserstrasse 12, 76131 Karlsruhe, Germany

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/01605682.2024.2436063>.

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

about the timing, type, and extent of a cyber-attack remain unknown beforehand.

In healthcare, these countermeasures are vital; any loss, modification, or restricted access to data can result in substandard care or even loss of life. While research on cybersecurity countermeasures has grown, it often overlooks the unique challenges of healthcare. With this research we make a first important contribution to address this gap by using stochastic programming, focusing on hospital functionality and prioritising the availability aspect of the CIA triad¹ to ensure uninterrupted medical care. Stochastic programming is particularly suited for addressing the uncertainties inherent in healthcare cybersecurity. It allows decision-makers to account for the unpredictable nature of cyber-attacks by enabling countermeasure implementations before attacks occur and adaptive responses post-incident. This flexibility is crucial in healthcare, where the consequences of disruptions can directly affect patient care. By considering various potential future attack scenarios, our stochastic approach provides robust solutions that mitigate risks, ensuring that the healthcare system remains resilient under diverse and uncertain conditions. Our numerical tests demonstrate the effectiveness of our approach in reducing patient harm, with the solution's robustness shown by its consistent performance across different parameter settings, emphasising the need for tailored cybersecurity strategies in healthcare.

This paper is organised as follows: the next section reviews the relevant literature on cybersecurity optimisation. A two-stage stochastic model for an NHS Trust is presented in Section 3, assuming a risk-neutral and risk-averse decision maker. Numerical tests are carried out in Section 4 and concluding remarks are given in Section 5.

2. Literature review of cybersecurity optimisation

Quantitative optimisation models can significantly support decision makers in the challenging task of improving cybersecurity. Due to the growing focus on optimisation models for improving cybersecurity, Karamdel et al. (2022) give a review in the field of cyber-physical power systems. The paper classifies optimisation models based on their applications and provides detailed insights into various techniques, including state estimation methods, game theory-based models, bi-level and multi-level optimisation frameworks, and heuristic methods for solving complex optimisation problems in cyber-physical systems.

One important issue in this context is the question of which countermeasures to invest to improve cybersecurity. Although many cybersecurity investment models

have been proposed recently, eg, Mai et al. (2021); Master et al. (2022); Mazzoccoli and Naldi (2022); Hyder and Govindarasu (2022); Gao et al. (2024), relevant literature in the healthcare context is scarce. One approach for choosing the optimal set of countermeasures is developed in Sönmez et al. (2022) where the objective is to minimise security risks in healthcare systems by applying mixed-integer linear programming. However, the uncertainty of cyber attacks is not taken into account, leading to the proposal of cybersecurity models that incorporate this factor (Bokhari et al., 2022). In healthcare, Attaallah et al. (2023) evaluate security risks using fuzzy logic-based techniques. Panda et al. (2020) use game theory and combinatorial optimisation to determine the optimal selection of cyber hygiene measures for healthcare staff, balancing the efficacy of different security measures against their costs.

Within stochastic programming, two-stage stochastic optimisation is frequently used in the cybersecurity domain. For instance, two-stage stochastic models are proposed in Novoa et al. (2018) to increase the cyber-resilience of Cyber-Physical Systems (CPS), to optimise staff scheduling of a 24/7 operating cybersecurity firm (Altner et al., 2018) and to enable cybersecure semiconductor wafer fabrication (Keith et al., 2024).

Papers more relevant in our context are Zhang et al. (2018), Paul and Zhang (2021), and Bhuiyan et al. (2021). The former develop a two-stage stochastic model with the objective of minimising costs for security control implementations and losses due to cyber-attacks. The authors model breach probabilities as a product of log-normal distributed random variables and convex piecewise linear functions. The two-stage stochastic model of Paul and Zhang (2021) optimises resource allocation by guiding firms on the mix of detection, prevention, and containment countermeasures, and advising government intelligence investments to minimise social costs and enhance detection effectiveness. Bhuiyan et al. (2021) use a bi-level stochastic network interdiction model to optimise cybersecurity resource allocation. The model uses a two-stage approach, where the first stage involves the defender deploying security measures on potential attack paths within a constrained budget, and the second stage models attackers' responses to these defences under uncertain budgets. The use of CVaR in the model aims to minimise both the expected maximum loss and the risk of significantly large losses from potential cyber-attacks, thereby improving the robustness of the cybersecurity strategy. A comparison table is available in the supplemental online material, showing the differences between our approach and other relevant methods. This table underscores the unique strengths of our model

within the context of healthcare cybersecurity. In the next section, we highlight the effectiveness of two-stage stochastic programming for cybersecurity tailored to the specific needs of the healthcare sector.

3. Cybersecurity stochastic programming

We begin this section by highlighting the advantages of stochastic programming for improving cybersecurity in healthcare in Section 3.1, followed by a short description of the NHS context in Section 3.2. In Section 3.3, we propose a cybersecurity two-stage stochastic model assuming a risk-neutral healthcare provider and in Section 3.4 we extend our model using the CVaR concept.

3.1. Two-stage stochastic programming for healthcare cybersecurity

The literature review reveals that optimisation models for improving cybersecurity in healthcare are scarce, often overlooking the uncertain nature of cyber attacks and failing to guide countermeasure selection. To the best of our knowledge, none of the existing stochastic cybersecurity models are tailored to the needs of the healthcare domain, where breaches can have severe consequences beyond financial losses, including patient harm, eg, from compromised medical devices. We propose a two-stage stochastic model for enhancing cybersecurity, specifically tailored to the unique requirements of healthcare providers. This stochastic approach is particularly well-suited for cybersecurity because it effectively addresses the inherent uncertainties and variability of cyber-attack scenarios. Given the diversity of threats, with significant variations in occurrence, type, and impact, such uncertainty demands a decision-making framework that can evaluate multiple potential future scenarios—a core strength of stochastic programming.

Two-stage stochastic models are crucial for decision-making under uncertainty. In the first stage, decisions like deploying intrusion-prevention technologies are made without full knowledge of an impending cyber-attack. Once an attack occurs and its impact is assessed, second-stage decisions focus on mitigation and system recovery. This approach enables a structured response to uncertainty, ensuring more effective and timely actions both before and after an attack.

A two-stage stochastic model enables healthcare providers to make proactive decisions despite incomplete information about threats and adjust responses following an attack. This proactive capability is crucial in preventing significant disruptions

to medical services, which directly impact patient care and safety.

The approach is also flexible, accommodating various cyber threats and countermeasures, allowing for sensitivity analysis and robust solutions. Tailorable to the specific budgets and operational needs of different healthcare organisations, it is a practical tool for real-world applications. This adaptability is vital in the rapidly evolving field of cybersecurity, helping healthcare organisations stay resilient against emerging threats.

Additionally, a two-stage stochastic model supports a systematic framework for evaluating the trade-offs between various cybersecurity investments and their expected outcomes. This approach ensures that resources are allocated efficiently, making cybersecurity investments both cost-effective and impactful. Unlike reactive strategies, proactive planning is essential for maintaining uninterrupted healthcare services.

The classic two-stage approach assumes a risk-neutral decision-maker, focusing on the most probable outcomes. However, rare but highly sophisticated attacks, though unlikely, can have devastating effects. A key advantage of two-stage stochastic models is their flexibility in incorporating different risk metrics.

In contrast, the *Value-at-Risk* (VaR) and *Conditional Value-at-Risk* (CVaR) have recently been used in optimisation (Filippi et al., 2020; Pavlikov et al., 2018), resulting in computationally tractable optimisation models. The objective with CVaR is to minimise the risk of high losses while incorporating the decision maker's risk preference. CVaR is particularly recommended for risk-averse decision makers due to its superior performance (Jaaman et al., 2011). For the first time, we present two-stage stochastic models to determine a set of countermeasures for a risk-neutral and a risk-averse decision maker, tailored to the needs of NHS Trusts.² While our research focuses on the NHS, healthcare providers worldwide share similar challenges and constraints, enabling our research to be generalisable to a broader audience.

In conclusion, a stochastic programming approach is well-suited for managing cybersecurity in healthcare due to its effectiveness in handling uncertainty, optimising risk, offering flexibility, and providing enhanced decision support. It delivers a scalable solution that addresses the critical needs of the healthcare industry.

3.2. NHS context

The WannaCry incident in 2017 was a spur for the NHS to improve its cyber resilience. For instance, all NHS Trusts had to obtain the Cyber Essentials Plus (CEP) certificate by June 2021. This scheme consists of basic countermeasures for information security,

which is in line with the CIA triad. Each NHS Trust is responsible for meeting the CEP requirements, which must be achieved with a limited budget and heavily scrutinised procurement processes.

The optimisation model we develop in this paper could help an NHS Trust decide which countermeasures to implement for prevention of cyber-attacks (first stage), to minimise its negative impact on patients in the aftermath (second stage). We focus particularly on cyber-attacks that compromise the availability of information and systems, leading to disruptions in medical care. The corresponding two-stage decision process is illustrated in Figure 1. The Trust can choose from a set of countermeasures with different effectiveness levels. These first-stage decisions have to be made under uncertainty as neither the attack type nor its outcome is known a priori. Uncertainties are captured by a set of scenarios containing details of possible cyber incidents and their impact on system and data availability. If a cyber-attack is successful and its scope on the Trust becomes evident, i.e. a specific attack scenario is realised, then second-stage decisions have to be taken. Depending on what systems and devices are affected or blocked, it may no longer be possible to provide some medical treatment. For instance, if clinical robotic surgical systems are hacked and become unusable, the Trust must decide how many patients can be treated manually and how many will need to be rescheduled, i.e., rejected. These decisions are made in the second stage, i.e. in the aftermath of an attack. The objective of our optimisation model is to minimise the expected number of manually treated or rejected patients over all possible cyber-attack scenarios. Choosing appropriate preparation measures before an attack occurs can significantly reduce its impact in the aftermath.

3.3. Risk-neutral cybersecurity two-stage stochastic model

The notation is as follows.

Sets:

I : Set of countermeasures (single countermeasures and bundles of countermeasures)

J : Set of mandatory single countermeasures, $J \subset I$

K : Set of treatment types

L : Set of intensity levels

S : Set of attack scenarios

Scenario-independent parameters:

B : Maximum available budget

C_k : Capacity, i.e., number of patients of type $k \in K$ that can be treated normally

c_{il} : Cost for countermeasure $i \in I$ at level $l \in L$

d_k : Demand, i.e. number of patients requiring treatment of type $k \in K$

γ_k : Importance weighting for treatment of type $k \in K$ with $\gamma_k \geq 0$ and $\sum_{k \in K} \gamma_k = 1$

κ : Importance weighting for rejected patients

U : Total manual capacity, i.e., maximum number of patients that can be treated manually

Scenario-dependent parameters:

C_k^s : Capacity reduction for treatment of type $k \in K$ in attack scenarios $s \in S$

e_{il}^s : Effectiveness of countermeasure $i \in I$ at level $l \in L$ on attack scenario $s \in S$ with $e_{il}^s \in (0, 1]$, i.e. $e_{il}^s = 1$ completely ineffective

m_i^s : 1, if countermeasure $i \in I$ has a mitigation effect on attack scenario $s \in S$, 0 otherwise

p^s : Probability of attack scenario $s \in S$

U_k^s : Number of patients of type $k \in K$ that can be treated manually in attack scenario $s \in S$

Scenario-independent decision variables (first stage):

x_{il} : 1, if countermeasure $i \in I$ is implemented at level $l \in L$; 0 otherwise

Scenario-dependent decision variables (second stage):

w_k^s : Number of manually treated patients of type $k \in K$ in scenario $s \in S$

y_k^s : Number of normally treated patients of type $k \in K$ in scenario $s \in S$

z_k^s : Number of rejected patients for treatment of type $k \in K$ in scenario $s \in S$

These parameters and variables are explained in more detail in the following using Figure 2 where some relationships are illustrated.




Figure 1. Two-stage decision process: first stage decisions on implementing preparation measures *before* and second-stage decisions on patient (non-)treatment *after* attack realisation.

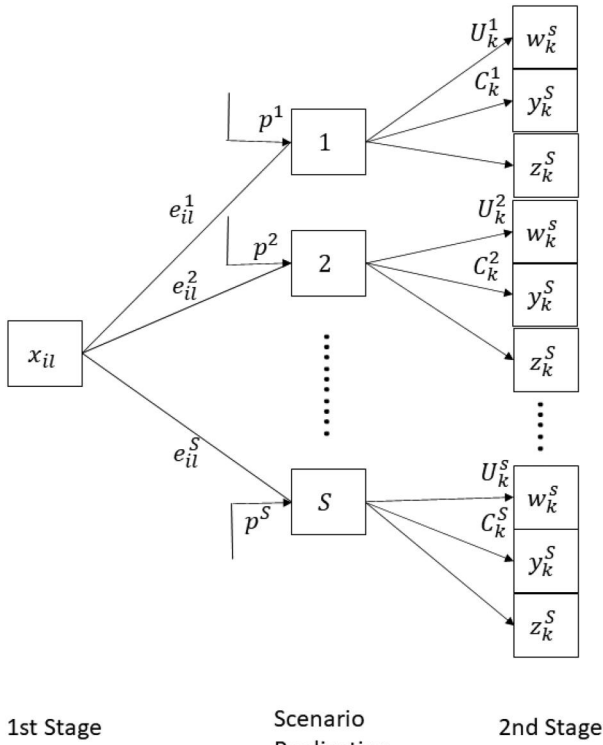


Figure 2. Impact of preventive countermeasure implementations x_{il} (first stage) on capacity in scenario s and patient treatment, i.e. number of normally treated y_k^s , manually treated w_k^s , and rejected patients z_k^s (second stage).

The Trust can reduce the impact of an attack by implementing countermeasure $i \in I$ at intensity level $l \in L$ before an attack occurs (first stage). For instance, employee training as a preparation measure can be performed annually ($l=1$), half-yearly ($l=2$) or quarterly ($l=3$).

In the following, we refer to a countermeasure i as a set of at least one single countermeasure, i.e. bundles. For instance, anti-virus software and firewall are often combined, protecting systems from malicious programs. If countermeasure i at level l is selected in the first stage, i.e. $x_{il} = 1$, then it can have a mitigation effect e_{il}^s on the impact of attack scenario s . Here, $e_{il}^s \rightarrow 0$ states that countermeasure i can almost completely reduce the impact of attack scenario s while $e_{il}^s = 1$ indicates that it is useless against s . As no countermeasure can completely block attacks in practice, $e_{il}^s \rightarrow 0$ but never reaches 0. A higher level l indicates a better effectiveness of countermeasure i , i.e. $e_{i1} > e_{i2} > \dots > e_{iL}$.

When the impact of a successful cyber incident becomes evident, i.e. an attack scenario has been realised, the infected hospital has to decide on how many patients requiring treatment of type k , i.e. d_k , can be treated normally despite the attack y_k^s , or by reverting to manual records w_k^s , or how many must be rescheduled or referred to other hospitals, i.e. rejected z_k^s , see Figure 2. These second-stage decisions depend on the outcome of attack scenario s . For instance, scenario s can represent a Patient

Administration System (PAS) hacking attack with a direct impact on the available capacity. Depending on the effectiveness level e_{il}^s of the implemented countermeasure x_{il} the maximum number of emergency patients that can be admitted through PAS is reduced by C_k^s . The remaining patients must be processed manually, but manual capacity U_k^s is limited and some patients might be rejected. This is because A&E staff are forced to work with pen and paper, which slows down processing times significantly. In the case of an A&E department closure all emergency patients must be turned away.

Choosing not to deploy preventive countermeasures in the first stage leads to more potential patient rejections in the second stage, as z_k^s strongly depend on pre-attack preparations. Anticipating the impact of future attacks on patients in the second stage can help to make better preparation decisions in the first stage.

The two-stage stochastic cybersecurity model is as follows:

Objective function

$$\min \sum_{s \in S} \sum_{k \in K} p^s \gamma_k (\kappa z_k^s + (1 - \kappa) w_k^s) \quad (1)$$

s.t.

$$\text{Maximum one intensity level allowed:} \quad (2)$$

$$\sum_{l \in L} x_{il} \leq 1 \quad \forall i \in I$$

Mandatory countermeasures must be implemented:

$$\sum_{l \in L} x_{il} \geq 1 \quad \forall i \in J \subset I \quad (3)$$

$$\text{Limited budget:} \quad \sum_{i \in I} \sum_{l \in L} c_{il} x_{il} \leq B \quad (4)$$

Limited capacity in the case that there is no or an ineffective counter measure:

$$y_k^s \leq C_k - C_k^s \left(1 - \sum_{i \in I} \sum_{l \in L} x_{il} m_i^s \right) \quad \forall k \in K, s \in S \quad (5)$$

Limited capacity in the case that an effective counter measure is in place:

$$y_k^s \leq C_k - C_k^s \left(\sum_{i \in I} \sum_{l \in L} m_i^s x_{il} e_{il}^s \right) \quad \forall k \in K, s \in S \quad (6)$$

Limited manual capacity:

$$w_k^s \leq U_k^s \quad \forall k \in K, s \in S \quad (7)$$

$$\sum_{k \in K} w_k^s \leq U \quad \forall s \in S \quad (8)$$

Number of rejected patients:

$$z_k^s = d_k - y_k^s - w_k^s \quad \forall k \in K, s \in S \quad (9)$$

Binary and non-negativity constraints:

$$x_{il} \in \{0, 1\} \quad \forall i \in I, l \in L \quad w_k^s, y_k^s, z_k^s \geq 0 \quad \forall k \in K, s \in S \quad (10)$$

The objective (1) is to minimise the expected number of rejected and manually treated patients

weighted by the scenario probability p^s and the importance factor γ_k . The latter indicates which type of treatment k is prioritised by the decision maker. The higher γ_k , the more important treatment k is and the more we want to minimise rejection. In addition, κ is chosen to give relatively more importance to rejections compared to manually processed patients.

Constraints (2) state that every countermeasure i can be implemented only at one level l , eg, staff training can take place either annually *or* quarterly.

As every NHS Trust has to adhere to internal and external requirements like CEP, constraints (3) state that the implementation of certain single countermeasures $J \subset I$ is mandatory. For instance, anti-malware software must be installed according to the CEP scheme.

Constraint (4) states that expenditure on countermeasure implementations cannot exceed the available budget B .

When an attack scenario realises and the corresponding impact on the capacity becomes evident, then second-stage decisions can be made. Not more patients can receive normal treatment k than the available capacity in the respective attack scenario s allows, see Constraints (5) and (6). If countermeasure i is not implemented, i.e. $x_{il} = 0$, although it could have a mitigating effect on attack scenario s , i.e. $m_k^s = 1$, the attack is assumed to reach its full potential and reduce the normal capacity C_k by C_k^s , see (5). The same holds if countermeasure i is implemented, i.e. $x_{il} = 1$, but is ineffective against attack scenario s , i.e. $m_k^s = 0$. In the case of $C_k = C_k^s$, no capacity is available to treat patients of type k normally so that $\gamma_k^s = 0$, eg, no patient can be admitted *via* PAS in the A&E unit.

In contrast, for $x_{il} = 1$ and $m_k^s = 1$ in (6), countermeasure i can attenuate the effect on capacity reduction C_k^s of attack scenario s by e_{il}^s . Constraints (6) ensure that only one countermeasure i is chosen and it will have the best effectiveness, i.e. when e_{il}^s is lower, for each attack scenario s . Note that it can be a different countermeasure for a different scenario. The formulation in (6) prevents additive effects of e_{il}^s by ensuring that implementing multiple countermeasures does not lead to a combined impact on the same cyber-attack, as only the most effective countermeasure is applied to each scenario.

Should the capacity not be sufficient in an attack scenario, some or all patients might be treated manually w_k^s , i.e. without technical equipment or digital system support. However, the maximum number of these patients depends on the manual capacity U_k^s in (7), which is limited by, eg, handling times. For instance, in an attack scenario where patient records are blocked, necessary information

or results have to be accessed by phone or fax limiting the hospital's processing capacities. This manual capacity depends on the treatment type k . Some attack scenarios may even result in $U_k^s = 0$, eg, when departments have to be closed, so that patients cannot be treated at all.

A total capacity U is available for treating patients manually, which limits the number of patients that medical staff can handle manually over all treatment types k , see (8).

According to (9), the number of rejected patients z_k^s is determined by the number of patients that can be treated normally y_k^s and manually w_k^s .

Constraints (10) define the binary nature of first-stage variables and non-negativity of second-stage variables.

Optimisation is performed over both stages by minimising the objective function in (1) *before* an attack occurs, taking all possible attack scenarios and second-stage decisions into account, to determine the best set of countermeasures.

3.4. Risk-averse cybersecurity two-stage stochastic model

In the previous model the decision maker relies on the expected objective value, i.e. scenarios with low probability but high impact have almost no influence on the decisions. However, worst-case scenarios should particularly be avoided within healthcare due to their catastrophic impact on patients. Therefore, we extend the previous model (Equations (1–10)) by incorporating the risk metric CVaR minimising the risk of rejecting a high number of patients. As explained in the supplemental online material, the confidence level α has to be chosen by the Trust to reflect its degree of risk aversion. Based on this choice, two additional types of decision variables have to be determined: η denotes the VaR and v^s is the number of rejected patients exceeding η in scenario s , i.e. $v^s = \max\{0, \sum_{k \in K} z_k^s - \eta\} \forall s \in S$. In other words, the Trust can be confident with probability α (eg, 95%) that the maximum number of patients who have to be rejected is η , and with the remaining probability $1 - \alpha$ the Trust has to reject more patients than η , denoted by v^s . These are the worst cases for which we would like to be prepared. In addition, the risk coefficient $\lambda \geq 0$ has to be chosen depending on the relative importance of CVaR for the respective Trust.

The risk-averse version of the previous model is as follows:

Bi-objective function

$$\min_{x, y, z, v, \eta} \sum_{s \in S} \sum_{k \in K} p^s \gamma_k (\kappa z_k^s + (1 - \kappa) w_k^s) + \lambda CVaR \quad (11)$$

s.t. (2) - (10) and
 Worst-case constraints

$$v^s \geq \sum_{k \in K} z_k^s - \eta \quad \forall s \in S \quad (12)$$

Non-negativity and real number specifications

$$v^s \geq 0 \quad \forall s \in S \quad \eta \in \mathbb{R}, \quad (13)$$

where $CVaR = \eta + \frac{1}{1-\alpha} \sum_{s \in S} p^s v^s$. Extending the objective function (1) by CVaR results in a bi-objective function (11) where not only the expected value over all scenarios but also the expected average outcome of worst-case scenarios is minimised. Constraint set (12) defines the number of rejected patients exceeding the VaR η for each scenario s . The last constraints (13) state that v^s is non-negative and η is a real number.

This risk-averse model allows the Trust to take necessary precautions not only for the most probable cases but also for the worst possible attack scenarios.

We give illustrative examples implementing these risk models in the next section to show differences in countermeasure decisions for a risk-neutral and risk-averse decision maker.

4. Computational experiments

In this section, we illustrate the implications of the risk-neutral and risk-averse optimisation model on a selected NHS Trust and its patients using an illustrative case study.

4.1. Illustrative case study

We designed a case study based on publicly available data and real data provided confidentially by an NHS Trust. In Sections 4.1.1 and 4.1.2, scenario-independent and scenario-dependent parameters are determined respectively.

4.1.1. Scenario-independent parameters

4.1.1.1. Countermeasures. For this example we assume seven types of countermeasures to primarily ensure the availability aspect of the CIA triad, namely, Anti-Malware (AM), Patch Management

(PM), Firewall and Secure Gateways (FW), Staff Training (ST), Access Control (AC), DDoS Prevention (DP) and User and Entity Behaviour Analytics (UEBA), see Table 1. Although many security controls exist (Federal Office for Information Security, 2023), expert discussions indicated that these seven countermeasures are considered the most important for ensuring availability. Each of them can be implemented at up to three levels of intensity l characterised by different effectiveness rates e_{il}^s . For instance, AM implemented at the first level can reduce the impact of a successful malware attack by 80% down to 20%, i.e. $e_{il}^s = 0.2$. The effectiveness rates are based on publicly available sources, but should be regarded as illustrative.

Countermeasure i refers to a single countermeasure (Table 1) or a bundle (Table 2) and can be effective against at least one attack type. For instance, the countermeasure bundle AM;FM can be used against malware and malicious traffic attacks, see Table 2. We follow the approach of Paul and Zhang (2021), determining e_{il}^s as the product of individual effectiveness rates. For instance, the effectiveness rate for countermeasure $i = AM;PM$ and $l = 1$ is the product of $e_{AM,1}^s = 0.2$ and $e_{PM,1}^s = 0.5$, namely $e_{AM;PM,1}^s = 0.1$.

Budget B and cost rates for implementing the described countermeasures c_{il} were provided by an NHS Trust.

4.1.1.2. Treatment types. All treatment provided by a hospital, either for an in- or outpatient, can be classified into four main categories:

- emergency care: patients treated at the A&E department
- diagnostics: digital imaging procedures, eg, X-ray radiography, tomography etc.
- surgery: all types of surgical interventions
- medication: all types of drug-based services

Based on NHS statistics for 2017/2018, we assume the following average numbers of patients requiring treatment k per day, i.e. d_k (Table 3). For the sake of simplicity, we set capacity C_k to the same values and do not prioritise between treatment

Table 1. Countermeasures with different effectiveness rates against attack types.

Countermeasure	Attack type	Effectiveness rate e_{il}^s			Source
		$l = 1$	$l = 2$	$l = 3$	
AM	Malware	0.2	0.1	0.02	Maimon (2019)
FW	Malicious traffic	0.4	0.1		Sophos (2018)
PM	Exploiting vulnerabilities	0.5	0.4	0.2	Gerace and Cavusoglu (2009)
AC	Unauthorised access	0.2	0.1	0.01	Marks (2019)
ST	Social engineering	0.3	0.15		Best (2019)
DP	DDoS attacks	0.06			Kotey et al. (2019)
UEBA	Malicious insiders	0.5	0.3		Logpoint (2020)

Table 2. Overview of countermeasure bundles effective against different attack types.

Countermeasure	Attack type					
	Malware	Malicious	Exploiting	Unauthorised	Phishing/SE	DDoS
Bundles		Traffic	Vulnerabilities	Access		
AM;FW	x	x				
AM;FW;PM	x	x	x			
AM;FW;PM;ST	x	x	x		x	
AC; ST				x	x	
DP;FW		X				x

Table 3. Daily average numbers of patients receiving treatment type k .

Treatment k	d_k (per day)
A&E attendances	809
Diagnostics	1340
Surgery	107
Medication	3082

Table 4. Possible attack types and the corresponding annual probabilities of occurrence and sources.

Attack type	Likelihood	Source
Malware	0.350	HIMSS (2019)
Malicious traffic	0.090	HIMSS (2019)
Exploiting vulnerabilities	0.570	BulletProof (2019)
Unauthorised access	0.050	HIMSS (2019)
Phishing/SE	0.600	HIMSS (2019)
DDoS	0.006	Verizon (2019)
Malicious insiders	0.060	HIMSS (2019)
Zero day	0.005	Assumption
Deepfake	0.001	Assumption

types, i.e. the criticality factor γ_k is the same for all k .

4.1.2. Scenario-dependent parameters

According to several annual statistics like HIMSS (2019), Verizon (2019), and Herjavec (2019), different types of cyber-attacks can occur with different probabilities, see Table 4. For instance, the probability that the initial point of compromise was a compromised medical or mobile device or a pre-loaded malicious software was 35% in 2018 (HIMSS, 2019).

We defined a total of 73 attack scenarios, see Table 5 for an extract,³ where the first scenario assumes no attack, i.e. there is no capacity reduction $C_k^s = 0$ and no need for manual capacities $U_k^s = 0$. Therefore, patient demand d_k can be satisfied under normal circumstances. In contrast, in the last nine scenarios the entire IT infrastructure and hence all departments may be affected. Except the worst cases, all other scenarios, i.e. attack types and consequences on various medical devices and systems, are derived from past cyber-attacks on healthcare such as WannaCry and can be considered plausible. To address the uncertainty of future threat impacts, the average numbers of patients in Table 3 were multiplied by a randomly generated percentage to determine the available capacity reduction C_k^s for treatment k in attack scenario s . For instance, in scenario 32 medical devices in the A&E department such as vital signs monitors are compromised by

unauthorised access, reducing capacity by 412 patients who now need to be monitored manually. The manual capacity U_k^s is determined based on the attack type and nurse-to-patient ratios (DropStat, 2014). Note that in some cases manual treatment is not possible. If MRI, CT scanners and X-rays are affected (scenario 42) there are no manual alternatives and therefore no manual capacity.

The attack likelihoods have to be scaled ensuring that the sum of daily probabilities over all scenarios is 1, i.e. $\sum_{s \in S} p^s = 1$.

4.2. Results

Both versions for the cybersecurity optimisation model introduced in Sections 3.3 and 3.4 were implemented in MATLAB 2020b and solved by MATLAB's build-in function `intlinprog`.

To show full use of model, the mandatory countermeasure implementation constraints in (3) are neglected for now. Based on discussions with NHS Trust authorities, worst-case and expected risks in the risk-averse optimisation model are seen equally critical. Therefore, we set $\lambda = 1$ in the objective function (11) so that the expected value as well as the CVaR are weighted equally. The first-stage solutions x_{ij} for the risk-neutral and risk-averse (with a confidence level of $\alpha = 0.99$) optimisation model are shown in Table 6.

According to Table 6, the solution for the risk-neutral approach is to implement four countermeasures in preparation for cyber-attacks, namely AM, FW, PM and ST. In contrast, the risk-averse model suggests those four plus two additional countermeasures, AC and UEBA. The implementation levels show that the risk-averse solution picks slightly less good protection against common lower impact risks in favour of a better defence against rare but potentially devastating risks like malicious access and insiders. Given budget constraints, it is impossible to implement all countermeasures, and the timing or nature of potential attacks is uncertain. Thus, even if not for DDoS attacks, these models provide a set of countermeasures that ensure readiness for common and critical scenarios, prioritising risk reduction for patient care continuity.

These different countermeasure selections lead to deviations in the model results, see Table 7. The

Table 5. Attack scenario s with daily probability of occurrence p^s , corresponding capacity reductions C_k^s and manual capacities U_k^s (excerpt).

Scenario s	Attack type	Scenario Probability p^s	Capacity reduction C_k^s for treatment type k				Capacity for manual processes U_k^s			
			Emergency	Diagnostics	Surgeries	Medication	Emergency	Diagnostics	Surgeries	Medication
1	No attack	6.33E-7	0	0	0	0	0	0	0	0
2	Malware	2.22E-2	0	0	35	0	0	0	54	0
12	Malicious traffic	5.70E-3	0	0	0	1553	0	0	0	1541
22	Exploiting vulnerabilities	3.61E-2	159	261	0	0	546	1005	0	0
32	Unauthorised access	3.17E-3	412	0	0	0	273	0	0	0
42	Phishing/social engineering	3.80E-2	0	819	0	0	0	0	0	0
47	DDoS	3.80E-4	0	0	7	0	0	0	54	0
57	Malicious insiders	3.80E-3	0	0	0	1856	0	0	0	1541
65	DDoS	5.43E-5	660	1079	86	2482	546	670	54	2312
66	Exploiting vulnerabilities	5.42E-3	669	1094	87	2517	546	670	54	2312
67	Phishing/social engineering	5.43E-3	698	1141	91	2625	546	670	27	1541
68	Malware	3.17E-3	721	1180	94	2714	273	670	27	1541
69	Malicious insiders	5.43E-3	757	1239	99	2848	273	335	27	1541
70	Malicious traffic	8.14E-4	785	1284	102	2953	273	335	0	771
71	Unauthorised access	4.52E-4	794	1299	104	2987	0	0	0	771
72	Zero day attack	3.17E-4	798	1306	104	3003	0	0	0	771
73	Deepfake attack	6.33E-5	813	1331	106	3061	0	0	0	771

Table 6. Comparison of countermeasure solutions x_{ij} with intensity levels l for the risk-neutral and -averse model.

Countermeasure x_{ij}	Intensity level l	
	Risk neutral	Risk averse
AM	2	1
FW	1	1
PM	3	2
ST	1	1
AC		1
UEBA		1

expected number of rejected patients as well as the objective value (sum of expected number of manually treated and rejected patients) for the risk-averse optimisation model are higher than for the risk-neutral approach. In the objective function (1) the most probable scenarios outweigh the worst cases through higher probabilities p^s . As the risk-neutral model focuses on minimising the impact of scenarios with highest probabilities rather than worst cases, the *expected* number of untreated patients and therefore the objective value is lower than in the risk-averse case. For instance, in the risk-neutral case, employing the highest patch management intensity reduces untreated patients during vulnerability exploitation attacks compared to the risk-averse case opting for intensity level 2 (Table 6). The crucial question revolves around the more effective approach in minimising high rejection rates. Comparing the risk-neutral and risk-averse optimisation models (1)–(10) and (11)–(13), respectively, the VaR and the CVaR, i.e. expected rejections in worst-case scenarios, are crucial. According to Table 7, $VaR = 1369$ indicates that with a probability of 99% not more than 1369 patients will be rejected. In the 1% of the remaining cases, an average of $CVaR = 1765$ untreated patients can be expected. Conversely, with 1% probability 2480 patients have to be rejected on average in the risk-neutral case, with $VaR = 1856$. This model might lead to over 5000 untreated patients in

Table 7. Comparison of the objective value, VaR, CVaR, expected, and worst-case number of rejections, and probability of rejections for the risk-neutral and -averse model.

	Risk neutral	Risk averse
Objective value	255	286
Expected rejections	119	122
VaR	1856	1369
CVaR	2480	1765
Worst-case rejections	5184	2901
Rejection probability	49%	37%

a worst-case scenario. In contrast, the risk-averse solution would cap rejections at 2900 even in the worst-case scenario. In addition, the risk of rejecting patients is higher in the risk-neutral case. The rejection probability in Table 7 measures the probability of positive second-stage variable z_k^s for at least one treatment type k and is defined as $\sum_{s=1}^S \{p^s : z_k^s = d_k - y_k^s - w_k^s > 0 \text{ for at least one } k \in K\}$, for a given first-stage solution. In nearly 50% of cases, the risk-neutral solution would result in patient rejections, which is 12% higher than the risk-averse solution.

The reason for the poorer performance of the risk-neutral approach is its focus on the most probable scenarios. According to Table 6, neither AC nor UEBA are part of the risk-neutral solution. As a result, worst-case attack scenario 69 (malicious insiders) and 71 (unauthorised access) in Table 5 cannot be fended off and would lead to numerous rejections in the case of realisation.

Note that if the optimisation model in Equations (1)–(10) is solved including constraint group (3), the same countermeasure selection as for the risk-averse approach is obtained.

The Value of the Stochastic Solution (VSS) measures the benefit of considering uncertainty in decision-making. It's calculated by comparing the outcome of a stochastic model (1)–(10) with that of a deterministic model, known as the expected value problem (EVP)(Birge & Louveaux, 2011). In this

example, the EVP has an objective value of 309, while the stochastic model's value is 255, resulting in a VSS of 54. This indicates that the deterministic model performs 21% worse in terms of manually treated and rejected patients. The VSS highlights the advantage of incorporating uncertainty, as it allows for optimal recourse actions after an attack, enabling dynamic adjustments like reallocating resources, rescheduling procedures, or transferring patients. This adaptability enhances the hospital's ability to respond effectively to unforeseen cyber threats, which deterministic models cannot achieve.

Most NHS Trusts are struggling to obtain the CEP certificate, let alone implementing additional non-mandatory countermeasures like staff training and UEBA. Consequently, preventing phishing/social engineering attacks and malicious insider threats becomes more difficult. Given the prevalence of phishing and social attacks (Table 4) and increase in deepfake threats, hospitals and patients might face significant risks without comprehensive staff training. Therefore, extending the CEP scheme by regular cybersecurity awareness courses would be reasonable.

We carried out sensitivity tests on the risk-averse optimisation model for budget B , risk parameter α , risk weight λ , scenario probabilities p^s and effectiveness levels e_{ij}^s that can be found in the supplement online material. Except for α , the results have revealed that the core selection of countermeasures x_{ij} remains unchanged, indicating its robustness. The analysis also shows that modest budget increases of 7% significantly improve countermeasure intensity and reduce the objective value, but further budget increases yield diminishing returns, making large investments less effective. For other NHS Trusts with different numbers of patients requiring different treatment types and a fundamentally different threat landscape, the optimal solution may lead to a different countermeasure selection.

Additional tests were conducted by simulating 1000 runs with randomly generated capacity reductions C_k^s to compare the quality of risk-neutral and risk-averse solutions, see the supplement online material. The results indicate that the risk-neutral solution leads to higher rejection rates and a greater number of rejected patients in worst-case scenarios, while the risk-averse solution significantly reduces these risks, offering better protection against high rejection probabilities.

5. Conclusion

We developed a stochastic optimisation model for the specific characteristics of the healthcare sector with the objective to minimise the negative impact

of cyber threats on patients by implementing a set of preparatory countermeasures. As highly sophisticated attacks may be rare but may have dramatic disruptive consequences for healthcare providers, we extended the optimisation model using the Conditional Value-at-Risk metric that enables risk-averse decision makers to be also prepared for worst-case scenarios. Numerical tests confirmed the advantages of the risk-averse model if focusing on the avoidance of rare but high-impact events. The corresponding solution is robust in the sense that neither changes to the risk weighting, budget, scenario probabilities nor effectiveness levels have altered the core selection of countermeasures. In particular, the solution suggests implementing an advanced analytics software such as UEBA as well as staff training in addition to the CEP requirements. Staff training is also chosen by the risk-neutral approach, i.e. independently of the decision maker's risk preferences this is an important countermeasure for improving information security and cyber resilience. As healthcare providers are already facing challenges in meeting the CEP guidelines, significant technical support, awareness at all levels, and an efficient budget allocation is required. The results of the case study have shown that even a moderate budget increase would lead to a considerable reduction of the expected number of rejected patients, improving the NHS Trust's security level.

In order to get a deeper understanding of the relative strengths and weaknesses of our approach, future research could focus on a detailed comparative analysis between our proposed model and other established cybersecurity strategies. Additionally, incorporating strategic interactions between attacker and defender could be achieved through a game-theoretic framework (Wu et al., 2022). This extension would enable our model to not only choose the most effective countermeasures but also anticipate and adapt to the strategies of cyber attackers, thus enhancing the realism and applicability of our approach.

Note that the scenarios in our case study consider single attacks, while multiple types of attacks might occur simultaneously. For instance, vulnerabilities can first be exploited to gain access to the system, followed by malware or ransomware attacks (Wixey, 2022). Future research could take this into account to create more realistic case studies.

The countermeasure solutions found by both optimisation models may be technically optimal but might not be accepted by hospital staff. For instance, multi-factor authentication is a powerful protection tool for access control but may represent a cumbersome and time-consuming hurdle in providing medical care, especially in emergency situations.

Therefore, a potential future research direction may be to incorporate human behaviour in order to find security solutions that are technically sophisticated but also accepted by medical workers.

Notes

1. CIA triad protects data through Confidentiality, Integrity, and Availability.
2. An NHS Trust comprises one or multiple hospitals.
3. The full table showing all 73 scenarios is given in the Supplemental Online Material.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Emilia Grass was supported by the German Fritz Thyssen Foundation under the Grant 4 O.18.O.O29WW

ORCID

Emilia Grass  <http://orcid.org/0000-0001-8460-8395>
 Christina Pagel  <http://orcid.org/0000-0002-2857-1628>
 Sonya Crowe  <http://orcid.org/0000-0003-1882-5476>

References

- Altner, D. S., Rojas, A. C., & Servi, L. D. (2018). A two-stage stochastic program for multi-shift, multi-analyst, workforce optimization with multiple on-call options. *Journal of Scheduling*, 21(5), 517–531. <https://doi.org/10.1007/s10951-017-0554-9>
- Attaallah, A., Al-Sulbi, K., Alasiry, A., Marzougui, M., Ansar, S. A., Agrawal, A., Ansari, M. T. J., & Khan, R. A. (2023). Fuzzy-based unified decision-making technique to evaluate security risks: A healthcare perspective. *Mathematics*, 11(11), 2554. <https://doi.org/10.3390/math11112554>
- Best, R. (2019). 35 UK cybersecurity statistics to make you serious about data protection. <https://www.infotech.co.uk/blog/35-cyber-security-stats-to-make-you-serious-about-data-protection>
- Bhuiyan, T. H., Medal, H. R., Nandi, A. K., & Halappanavar, M. (2021). Risk-averse bi-level stochastic network interdiction model for cyber-security risk management. *International Journal of Critical Infrastructure Protection*, 32, 100408. <https://doi.org/10.1016/j.ijcip.2021.100408>
- Birge, J., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer.
- Bokhari, S., Hamrioui, S., & Aider, M. (2022). Cybersecurity strategy under uncertainties for an IoE environment. *Journal of Network and Computer Applications*, 205, 103426. <https://doi.org/10.1016/j.jnca.2022.103426>
- BulletProof. (2019). *Annual cyber security report*. <https://www.bulletproof.co.uk/industry-reports/2019.pdf>
- DropStat. (2014). *Safe staffing ratios are crucial to reducing patient mortality rates*. <https://dropstat.com/blog/safe-staffing-nursing/>
- Federal Office for Information Security. (2023). *IT-Grundschatz-Compendium*. https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Standards-und-Zertifizierung/IT-Grundschatz/it-grundschatz_node.html
- Filippi, C., Guastaroba, G., & Speranza, M. G. (2020). Conditional value-at-risk beyond finance: A survey. *International Transactions in Operational Research*, 27(3), 1277–1319. <https://doi.org/10.1111/itor.12726>
- Gao, X., Gong, S., Wang, Y., & Zhang, Y. (2024). Information sharing and security investment for substitutable firms: A game-theoretic analysis. *Journal of the Operational Research Society*, 75(4), 799–820. <https://doi.org/10.1080/01605682.2023.2210594>
- Gerace, T., & Cavusoglu, H. (2009). The critical elements of the patch management process. *Communications of the ACM*, 52(8), 117–121. <https://doi.org/10.1145/1536616.1536646>
- Herjavec. (2019). *The 2020 healthcare cybersecurity report by cybersecurity ventures*. <https://www.herjavecgroup.com/2020-healthcare-cybersecurity-report-cybersecurity-ventures/>
- HIMSS. (2019). *Cybersecurity survey final report healthcare*. <https://www.himss.org/himss-cybersecurity-survey>
- Hyder, B., & Govindarasu, M. (2022). A novel methodology for cybersecurity investment optimization in smart grids using attack-defense trees and game theory. In: Mohammad Shahidehpour (Ed.), *2022 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)* (pp. 1–5). IEEE. <https://doi.org/10.1109/ISGT50606.2022.9817467>
- ITPro. (2024). *Healthcare sector cyber attacks are surging at an alarming rate, prompting frantic alerts by the FBI and CISA*. <https://www.itpro.com/security/healthcare-sector-cyber-attacks-are-surging-at-an-alarming-rate-prompting-frantic-alerts-by-the-fbi-and-cisa>
- Jaaman, S. H., Lam, W. H., & Isa, Z. (2011). Different downside risk approaches in portfolio optimisation. *Journal of Quality Measurement and Analysis JQMA*, 7(1), 77–84.
- Karamdel, S., Liang, X., Faried, S. O., & Mitolo, M. (2022). Optimization models in cyber-physical power systems: A review. *IEEE Access*, 10, 130469–130486. <https://doi.org/10.1109/ACCESS.2022.3229626>
- Keith, K., Castillo-Villar, K. K., & Bhuiyan, T. H. (2024). Attack graph-based stochastic modeling approach for enabling cybersecure semiconductor wafer fabrication. *Computers & Industrial Engineering*, 188, 109912. <https://doi.org/10.1016/j.cie.2024.109912>
- Kotey, S. D., Tchao, E. T., & Gadze, J. D. (2019). On distributed denial of service current defense schemes. *Technologies*, 7(1), 19. <https://doi.org/10.3390/technologies7010019>
- Logpoint. (2020). *What is user and entity behavior analytics? A complete guide to UEBA, how it works, and its benefits*. <https://www.logpoint.com/en/blog/ueba-user-and-entity-behavior-analytics/>
- Mai, V. S., La, R. J., & Battou, A. (2021). Optimal cybersecurity investments in large networks using sis model: Algorithm design. *IEEE/ACM Transactions on Networking*, 29(6), 2453–2466. <https://doi.org/10.1109/TNET.2021.3091856>
- Maimon, D. (2019). *Existing evidence for the effectiveness of antivirus in preventing cyber crime incidents*. <https://>

- scholarworks.gsu.edu/cgi/viewcontent.cgi?article=1000&context=ebscs_tools
- Marks, G. (2019). *Microsoft: Multi-factor authentication is 99 percent effective*. <https://www.forbes.com/sites/quick-erbettertech/2019/09/01/microsoft-multi-factor-authentication-is-99-percent-effectiveand-other-small-business-tech-news-this-week/>
- Master, A., Hamilton, G., & Dietz, J. E. (2022). Optimizing cybersecurity budgets with attack simulation. In *2022 IEEE International Symposium on Technologies for Homeland Security (HST)* (pp. 1–7). IEEE. <https://doi.org/10.1109/HST56032.2022.10024984>
- Mazzoccoli, A., & Naldi, M. (2022). Optimizing cybersecurity investments over time. *Algorithms*, 15(6), 211. <https://doi.org/10.3390/a15060211>
- Novoa, C., Siddique, K., Guirguis, M., & Tahsini, A. (2018). A game-theoretic two-stage stochastic programming model to protect CPS against attacks. In *INDIN* (pp. 15–22). IEEE.
- Panda, S., Panaousis, E., Loukas, G., & Laoudias, C. (2020). Optimizing investments in cyber hygiene for protecting healthcare users. In: Alessandra Di Pierro, Pasquale Malacaria, Rajagopal Nagarajan (Eds.), *From Lambda calculus to cybersecurity through program analysis: Essays dedicated to Chris Hankin on the occasion of his retirement* (pp. 268–291). Springer.
- Paul, J., & Zhang, M. (2021). Decision support model for cybersecurity risk planning: A two-stage stochastic programming framework featuring firms, government, and attacker. *European Journal of Operational Research*, 291(1), 349–364. <https://doi.org/10.1016/j.ejor.2020.09.013>
- Pavlikov, K., Veremyev, A., & Pasiliao, E. L. (2018). Optimization of value-at-risk: Computational aspects of MIP formulations. *Journal of the Operational Research Society*, 69(5), 676–690. <https://doi.org/10.1057/s41274-017-0197-4>
- Sönmez, F. Ö., Hankin, C., & Malacaria, P. (2022). Decision support for healthcare cyber security. *Computers & Security*, 122, 102865. <https://doi.org/10.1016/j.cose.2022.102865>
- Sophos. (2018). *The problem with next-gen firewall protection*. <https://news.sophos.com/en-us/2018/01/01/the-problem-with-next-gen-firewall-protection/>
- Verizon. (2019). *2019 Data breach investigations report*. <https://enterprise.verizon.com/resources/reports/dbir/>
- Wixey, M. (2022). *Multiple attackers: A clear and present danger. A sophos X-ops active adversary whitepaper*. <https://www.securityweek.com/cyberattack-victims-often-attacked-multiple-adversaries-research/>
- Wu, Y., Xiao, H., Dai, T., & Cheng, D. (2022). A game-theoretical model of firm security reactions responding to a strategic hacker in a competitive industry. *Journal of the Operational Research Society*, 73(4), 716–740. <https://doi.org/10.1080/01605682.2020.1854631>
- Zhang, H., Chari, K., & Agrawal, M. (2018). Decision support for the optimal allocation of security controls. *Decision Support Systems*, 115, 92–104. <https://doi.org/10.1016/j.dss.2018.10.001>