

Future Directions in Stochastic Cyber Risk Management*

Jerzy Dorobisz

Affiliation e.g. Institute of Information Technology and Cyber-security,
Faculty of Cybernetics, WAT, 2 Gen. Sylwestra Kaliskiego St., 00-908 Warsaw

Correspondence should be addressed to: Jerzy Dorobisz, jerzy.dorobisz@wat.edu.pl

* Presented at the 46th IBIMA International Conference, 26-27 November 2025, Ronda, Spain

Abstract

This study addresses the critical need for advanced quantitative approaches in cybersecurity risk management, motivated by the limitations of traditional qualitative methods in today's dynamic threat landscape. The research identifies a significant void in literature concerning the integration of sophisticated stochastic modeling techniques with practical enterprise risk management frameworks. Using a comprehensive methodological approach that combines systematic literature review and conceptual analysis, the paper examines five key development directions in stochastic cyber risk management. The findings reveal that while significant advancements are being made in machine learning, agent-based simulations, and extreme value theory, substantial challenges remain in data quality, computational scalability, and model interpretability. The study concludes that successful implementation requires interdisciplinary collaboration and that parametric insurance and risk securitization represent particularly promising avenues for transferring cyber risks to capital markets, though these require further methodological refinement.

Keywords: stochastic cyber risk management, enterprise risk management, machine learning, parametric insurance, risk quantification, cyber catastrophe bonds, digital twins

Introduction

The future of stochastic cyber risk management is moving towards integrated, dynamic models based on rich data.

Key development directions:

1. **Integration with enterprise risk management (ERM):** Cyber models will be incorporated into a company's overall risk portfolio, measured by common metrics (e.g., VaR), to show correlations with operational, financial, and reputational risk.
2. **Blockchain-enabled data consortia:** Anonymous, peer-to-peer sharing of incident data through decentralised ledgers will create the rich data sets necessary for accurate probability modelling.
3. **Dynamic AI models and "digital twins":** Artificial intelligence and machine learning will continuously update model parameters based on threat intelligence feeds. Simulations of attacks on faithful virtual replicas of infrastructure (digital twins) will provide insight into cause-and-effect relationships and potential financial impact.

4. **Parametric insurance and securitisation:** Parametric insurance, where payment is made after an objective condition is met (e.g., X hours of downtime), will develop. Risk will also be transferred to capital markets through financial instruments such as cyber catastrophe bonds (cat bonds).
5. **Quantification of systemic risk:** Advanced stochastic models will be developed to examine cascading effects and interdependencies in the network in order to measure the risk to the entire system (e.g., cloud provider failure).

Introduction: context and stochastic imperative

Today's cyberspace is an environment of unprecedented dynamism and complexity. The increasing frequency, sophistication and global reach of cyber attacks exceed the capabilities of traditional risk management methods (World Economic Forum, 2023).

Deterministic risk matrices, checklists and compliance-only approaches are proving increasingly inadequate in the face of fundamental uncertainty, high volatility and complex, non-linear interactions that characterise IT systems and their environment.

Qualitative approaches are particularly criticised, as they often fail to provide a solid basis for quantitative comparison of different types of risk, effective allocation of limited security resources, or economic justification for investments in protective measures.

In this context, stochastic cyber risk management emerges as a necessary evolution. It is defined as a systematic process of identifying, analysing, assessing, mitigating and monitoring cyber risk, based on the use of probabilistic models to represent and quantify uncertainty about key aspects of risk (Alfredsson & Verendel, 2017).

The basic principle of stochastic cyber risk management is the recognition of the randomness of cyber events and the fundamental lack of complete knowledge about all factors affecting system security.

Risk Typologies in ERM Context

Understanding various risk typologies is essential for effective ERM integration. Cyber risks can be categorized as:

Technical risks: Vulnerabilities in systems and networks

Operational risks: Process failures and human errors

Strategic risks: Poor decision-making regarding cybersecurity investments

Compliance risks: Regulatory violations and legal liabilities

Reputational risks: Damage to brand and stakeholder trust (NIST, 2018)

The mathematical foundations of stochastic cyber risk management include probability theory, stochastic processes, statistics, and risk calculation methodologies (Alfredsson & Verendel, 2017).

Current challenges further motivate the development and implementation of stochastic cyber risk management. The lack of historical data is particularly acute for new types of attacks and extreme events, where the so-called "long tails" of distributions remain virtually unknown (Coles, 2001).

The complexity of systems makes it difficult to fully model all dependencies and potential points of failure. The rapidly changing technological environment and evolution of adversarial tactics mean that risk models quickly become outdated (ENISA, 2023).

Main directions of development of stochastic risk management methods and technologies

Advanced machine learning (ML) and deep learning (DL) as the basis for stochastic modelling

Advanced machine learning techniques are becoming a fundamental pillar of modern stochastic cyber risk management, offering powerful tools for dealing with its complexity (Apruzzese et al., 2023). A key area of application is the discovery of risk distributions. Traditional models often rely on assumptions of simple

distributions which may not reflect the actual complex nature of cybersecurity data. ML and DL, in particular models such as Generative Adversarial Networks (GANs) allow for the discovery of nonlinear, multidimensional probability distributions without restrictive a priori assumptions (Goodfellow et al., 2016). Another groundbreaking application is probabilistic event prediction. Traditional predictive models often provide single, point forecasts (Bridges et al., 2019). However, in risk management, understanding the full probability distribution of possible future outcomes is crucial. ML and DL models, especially neural networks with probabilistic layers are capable of predicting not only the probability of an attack, but also the full frequency and severity distributions of attacks over specific time horizons. An important area is stochastic anomaly and threat detection. Conventional detection systems often operate in binary mode, generating an avalanche of false positives that overload SOC teams. The stochastic approach is based on models assessing the probability that a given observed anomaly is a real threat. Techniques such as Variational Autoencoders (VAEs) are particularly useful because they learn to model the distribution of data representing "normal" system behavior (Hindy et al., 2020).

Parametric Insurance and Securitization: Detailed Analysis

Methodological Foundations

Parametric insurance represents a significant evolution in cyber risk transfer, moving away from traditional indemnity-based policies toward objective, trigger-based mechanisms (Eling & Wirfs, 2019). The methodology involves:

1. **Trigger Definition:** Establishing measurable parameters that correlate strongly with cyber incidents
2. **Model Validation:** Ensuring statistical significance between triggers and actual losses
3. **Pricing Models:** Developing stochastic models that accurately price the risk

Analysis Framework

The relationship between methodology, analysis, and this development direction can be structured as follows:

Data Collection & Processing

- Gathering historical incident data with precise timing and impact metrics
- Correlating technical parameters (downtime, data records affected) with financial impacts
- Establishing baseline metrics for normal operations

Model Development

- Creating stochastic models that simulate various cyber incident scenarios
- Testing trigger sensitivity and specificity across different incident types
- Validating models against historical data where available

Risk Transfer Mechanism Design

- Structuring insurance products with clear parametric triggers
- Developing catastrophe bond frameworks for capital market risk transfer
- Establishing claims validation protocols

Practical Implementation Considerations

For organizations considering parametric insurance, several factors must be addressed:

Advantages:

- Rapid claims payment without lengthy loss adjustment
- Reduced basis risk through carefully designed triggers
- Complement to traditional insurance coverage

Challenges:

- Basis risk - the potential mismatch between actual loss and trigger parameters
- Model risk - inaccuracies in the underlying stochastic models
- Data quality - requiring robust monitoring and measurement systems

Future Applications

The potential applications of parametric cyber insurance extend beyond traditional coverage:

- **Supply Chain Resilience:** Parametric triggers for supplier cyber incidents
- **Systemic Risk Transfer:** Catastrophe bonds for critical infrastructure protection
- **Real-time Risk Management:** Dynamic premium adjustments based on continuous risk assessment

This development direction bridges the gap between technical risk assessment and financial risk transfer, creating new opportunities for organizational resilience.

difficult to access, incomplete or error-prone. There is a significant risk of systematic errors in training data, which can lead to unfair or inaccurate predictions. ML models are also vulnerable to deliberate adversarial attacks, such as data poisoning or the generation of specially crafted examples (*adversarial examples*) designed to deceive the model. Model validation, especially in the context of rare, high-impact events, remains difficult. Furthermore, the computational costs associated with training large deep models and performing real-time inference are significant.

Agent-based stochastic simulations (ABSS)

The development and practical implementation of advanced methods faces serious, interrelated challenges that require an interdisciplinary approach to solve.

Data challenges are often a fundamental barrier. Data quality and completeness are crucial, but in practice, cybersecurity data is often contaminated and lacks format standardisation (Zuech et al., 2015).

The amount and representativeness of data, especially for rare extreme events is usually insufficient. There is also a problem of lack of representativeness – the data available for model training may not cover all types of environments.

Computational and scalability challenges are inherent in the complexity of future methods. The demand for computing power is enormous. Training deep neural models on large security log datasets requires powerful High-Performance Computing solutions.

Interpretability, trust and acceptance are crucial for practical implementation. The "black box" problem refers to a situation where even the model creators find it difficult to understand how complex ML models arrived at a particular prediction (Goodman & Flaxman, 2017).

The solution to this dilemma is the intensive development of Explainable AI (XAI). Methods such as LIME and SHAP are being adapted for cyber risk models.

Agent-Based Stochastic Simulations (ABSS) offer a powerful framework for modelling the most complex aspects of cyber risk, where traditional analytical methods fail. The essence of ABSS is to simulate the behaviour of heterogeneous, autonomous "agents" and their random interactions in a virtual environment that mirrors cyberspace. Agents represent various entities: end users with varying levels of awareness and behaviour, attackers with different motivations (financial, espionage, activism), skills and resources, defence systems (firewalls, IDS/IPS systems, EDR), as well as infrastructure elements (network nodes, servers, IoT devices). Each agent has a defined set of attributes (e.g., privilege level, list of known vulnerabilities, level of "susceptibility" to phishing) and behaviour rules, often expressed probabilistically (e.g., probability of clicking on a phishing link, probability of scanning the network for vulnerabilities by an attacker). Some agents may also have the ability to learn and adapt during the simulation. A key advantage of ABSS is its ability to capture emergent phenomena – complex patterns and behaviours of the entire system that arise from simple interactions between individual agents and would be impossible to predict based on an analysis of individual components. This allows for the study of cascading effects (when the failure of one element leads to the failure of others) and non-linear feedback loops in dynamically changing conditions.

ABSS is particularly valuable for testing "What-If Analysis" scenarios. It allows for virtual experimentation with various changes in the environment without risk to actual systems. It is possible to simulate the impact of changes infrastructure configuration, the implementation of new security controls, the emergence of a new type of threat (e.g., an exploit for a critical zero-day vulnerability), or a change in an adversary's strategy (e.g., a shift from mass attacks to targeted attacks). By running multiple simulations for a given scenario (e.g., using the Monte Carlo method, where input parameters are randomly selected from appropriate probability distributions), it is possible to obtain not a single answer, but a complete probability distribution of different outcomes (e.g., distribution of time to full system compromise, distribution of expected number of infected hosts, distribution of expected financial losses). This probabilistic approach provides much richer and more useful information for decision-makers than deterministic simulations.

Key applications of ABSS in stochastic cyber risk management include:

1. Modelling the spread of malware: simulating how a virus, worm or ransomware spreads within a corporate network or between different organisations, taking into account network topology, the effectiveness of existing security measures, user behaviour and attacker actions.
2. Assessing the effectiveness of phishing campaigns: modelling how different user profiles respond to different types of phishing messages, and how awareness-raising activities (training, simulated phishing attacks) reduce the success rate of attacks.
3. Network resilience analysis: testing how the network responds to DDoS attacks, physical damage to key nodes, or coordinated targeted attacks. Simulations allow you to assess which infrastructure elements are most critical and where investments in redundancy or more resilient protocols will yield the greatest benefit.
4. Supply chain risk modelling: simulating how a security incident at one supplier (e.g., compromised software update) propagates stochastically to the security of the target organisation and its other partners. This allows weak links in the ecosystem to be identified.
5. Testing incident response strategies (*Incident Response Playbooks*): verifying the effectiveness and efficiency of developed response procedures for various incident scenarios in a controlled but realistic simulation environment.

Despite its potential, ABSS faces serious challenges. The extremely high complexity of modelling requires interdisciplinary knowledge combining cybersecurity, stochastic modelling, sociology (for user behaviour) and game theory (for attacker-defender interactions). The difficulties in calibrating and validating agent models are enormous – how can one reliably determine the parameters describing the behaviour of an attacker agent or the probability of a user making a risky decision? Extreme computational requirements are another barrier; simulations of large, complex environments with thousands or millions of agents can require powerful computing clusters and take hours or days to complete. Obtaining reliable data to define agent rules and parameters, especially for attacker behaviour (which is often hidden), remains a significant problem.

Interdisciplinary challenges and considerations

Ethical, legal and regulatory issues are becoming increasingly important as the capabilities of stochastic cyber risk management systems grow.

Privacy is threatened by the intensive collection of employee behavioral data necessary for training effective models. Compliance with stringent regulations, such as GDPR, is mandatory and poses challenges for system designers.

The bias and fairness of algorithms is another pressing issue. If models learn from historical data that reflects existing biases, they may perpetuate these inequalities (Barocas et al., 2019).

Responsibility becomes unclear when an autonomous system makes a wrong decision that causes damage. Similarly, responsibility for errors in a stochastic model leading to misallocation of security resources requires legal regulation (Veale et al., 2018).

Conclusion and Future Outlook

The development and practical implementation of advanced methods of stochastic cyber risk management faces a number of serious, interrelated challenges that require an interdisciplinary approach to solve.

Data challenges are often a fundamental barrier. Data quality and completeness are crucial, but in practice, cybersecurity data is often contaminated (e.g., erroneous logs), lacks format standardisation (making aggregation difficult), omits key contextual attributes, and sometimes deliberately introduces misinformation. The amount and representativeness of data, especially for rare extreme events (the so-called "long tails" of distributions) or completely new threats, is usually insufficient. There is also a problem of lack of representativeness – the data available for model training may not cover all types of environments, sectors or scenarios, leading to models that only work well in narrow contexts. The availability and sharing of security incident data is severely limited due to organisations' concerns about reputation, legal liability and the disclosure of sensitive information about vulnerabilities. Barriers to sharing threat intelligence between organisations, and even within sectors, make it difficult to build comprehensive risk pictures. Ethical issues surrounding data collection are also pressing: monitoring detailed user behaviour

to model the human factor raises serious privacy concerns; the use of data from the Dark Web raises questions about the legality and ethics of its acquisition; finally, there is a real risk that ML algorithms, trained on biased data, will perpetuate or even reinforce existing social biases, leading to discriminatory results.

Computational and scalability challenges are inherent in the complexity of future methods of stochastic cyber risk management. The demand for computing power is enormous. Training deep neural models on large security log datasets, conducting detailed ABSS simulations of large corporate environments or critical infrastructure, or running multiple Monte Carlo simulations for complex risk models with thousands of parameters – all these tasks require powerful High-Performance Computing (HPC) and/or scalable cloud solutions. The execution time of some simulations, especially very detailed ones (such as high-fidelity ABSS), can be measured in hours or days. This makes it impossible to use them to support real-time decisions or in rapidly changing crisis situations. The costs associated with such computing infrastructure, specialised software and the employment of highly qualified personnel (data scientists, ML engineers, simulation specialists) are very high, which may limit the availability of these methods to smaller organisations. Therefore, algorithm optimisation is of key importance. There is a need to develop more efficient approximation algorithms that provide good approximations at lower costs, dimensionality reduction techniques that allow working on relevant features without losing key information, and methods that accelerate convergence in the machine learning and RL process.

Interpretability, trust and acceptance are crucial for practical implementation, especially in the case of deep learning-based methods. The "black box" problem refers to a situation where even the model creators find it difficult to understand how complex ML/DL/RL models (especially deep neural networks with millions of parameters) arrived at a particular prediction, recommendation or decision. The inability to logically explain and verify the model's decision-making process is a serious drawback. This leads to a lack of trust among decision-makers – senior management, internal and external auditors, as well as regulatory authorities.

These individuals may prefer traditional, less accurate but more transparent methods, as they allow them to understand the rationale and take responsibility. The solution to this dilemma is the intensive development of the field of Explainable AI (XAI). Methods such as LIME (Local Interpretable Model-agnostic Explanations), SHAP (SHapley Additive exPlanations) and counterfactual explanations ("what would have to change for the result to be different?") are being adapted and developed specifically for the specifics of cyber risk models. However, the main challenge for XAI in stochastic cyber risk management is to find a way to explain probabilistic results and complex relationships between multiple risk variables in a way that is understandable and useful to non-technical stakeholders making strategic decisions.

Adversarial Machine Learning poses a direct threat to the reliability and effectiveness of stochastic cyber risk management systems. Stochastic models themselves, especially those based on ML, can become targets of attacks:

- **Data Poisoning:** Attackers can deliberately introduce erroneous or manipulated data into the model's training set. The goal may be, for example, to underestimate the risk assessment for specific systems or types of attacks so that they remain undetected or low-priority, or, conversely, to overestimate the risk in order to cause misinformation and inefficient allocation of defence resources.
- **Adversarial Attacks:** These involve deliberate, subtle modification of input data (e.g., features describing a file sample, network packet, or user behaviour) in real time (during model inference). These modifications are often imperceptible to humans, but cause the model to classify a malicious object as safe or vice versa.
- **RL Environment Attacks:** If an RL agent is trained in a simulated environment, an attacker may attempt to manipulate that environment or the sensory data provided to the agent in order to teach it suboptimal or even harmful behaviours (e.g., disabling key security measures under certain conditions).
- **Model Stealing/Extraction:** By interacting with the model's API (e.g., sending queries and analysing responses), an attacker may attempt to reconstruct its internal structure or decision-making function, enabling better preparation of attacks that bypass its defence mechanisms.

Building robustness of models against such attacks is therefore becoming a priority. It is necessary to design models with security in mind from the outset (Security by Design), use data manipulation detection techniques (data sanitisation, anomaly detection on input data), implement mechanisms for continuous validation of models on new, verifiable data, and closely monitor their behaviour for unexpected anomalies or a decline in effectiveness.

Ethical, legal and regulatory issues are becoming increasingly important as the capabilities and autonomy of stochastic cyber risk management systems grow. Privacy is threatened by the intensive collection of employee behavioural data, telemetry data from end devices, and detailed network logs necessary for training effective models. Compliance with stringent regulations, such as the EU's GDPR and California's CCPA, is mandatory and poses a challenge for system designers. The bias and fairness of algorithms is another pressing issue. If stochastic cyber risk management models learn from historical data that reflects existing social or organisational biases (e.g. regarding certain departments or user groups), they may perpetuate or even reinforce these inequalities. An example would be an insider threat assessment model that unfairly targets employees from certain demographic groups or favours protecting some parts of the organisation at the expense of others due to biased training data. Regular audits of algorithms for fairness and the implementation of techniques to ensure fairness are necessary. Responsibility becomes unclear when an autonomous RL-based system makes a wrong decision (e.g., incorrectly blocking a critical service or ignoring a real attack) that causes damage. Who is responsible: the model developers, the software provider, the end user implementing the system, or perhaps the "artificial intelligence" itself? Similarly, responsibility for errors in a stochastic model leading to misallocation of security resources and resulting in an incident requires legal regulation. Transparency and supervisability are key requirements, especially in light of developing regulations such as the EU AI Act. How can we ensure an adequate level of transparency in the operation of "black box" models without revealing sensitive information about their operation to potential attackers? How can meaningful human oversight (Human-in-the-Loop) of critical decisions made by autonomous defence systems? Finally, national and social security is called into question in the context of the use of advanced stochastic cyber risk management models by state actors for surveillance or cyber warfare operations. There is also an ethical dilemma associated with the development of advanced adversarial models – while they serve to better prepare defences, they can also be captured or replicated by attackers to improve their own attacks.

The gap between theory and practice remains significant. The lack of qualified personnel combining in-depth knowledge of cybersecurity, statistics, machine learning and data engineering is a serious limitation to the implementation of advanced methods of stochastic cyber risk management in organisations. The complexity of implementing these often highly sophisticated methods into existing business processes, security tools and organisational culture is enormous, especially in companies with lower levels of cybersecurity maturity, where basic practices may not yet be well established. Finally, the issue of cost and return on investment (ROI) is difficult to resolve unequivocally. Demonstrating a clear, quantitative return on investment in complex and costly stochastic cyber risk management solutions, especially in the short term, is a challenge for security departments seeking budgets. The benefits in terms of avoided losses or increased resilience are often difficult to measure precisely and demonstrate *ex ante*.

Summary, conclusions and outlook

The future of stochastic cyber risk management is inextricably linked to accelerating technological progress, particularly in artificial intelligence and machine learning.

As demonstrated in this paper, key development directions focus on deeper integration of AI/ML into modeling complex risk distributions and predicting events under conditions of uncertainty.

However, the effective implementation of these promising directions requires absolute interdisciplinarity. The future lies in close cooperation between specialists in cybersecurity, mathematics, data engineering, and domain-specific areas.

Ethics and responsibility must be the foundation for development. The principles of transparency, fairness, and meaningful human oversight must be built into systems from the outset.

The conclusion is clear: stochastic cyber risk management is no longer an academic curiosity, but has become a necessary foundation for building truly resilient organizations. Investment in further research and practical implementation is crucial for security and economic stability.

However, the effective implementation of these promising directions requires absolute interdisciplinarity. The

future of stochastic cyber risk management lies in close cooperation between specialists in the fields of cybersecurity, mathematics and stochastics, data engineering and machine learning, psychology and sociology (behavioural modelling), law, ethics and specific domain areas (e.g. energy, finance, healthcare). Only teams combining this diverse knowledge are able to meet the complexity of the challenges facing quantitative cyber risk management.

Adaptability and automation are becoming imperative. The speed at which the cyber threat landscape and technology are changing requires stochastic cyber risk management systems to be highly adaptive – capable of continuous learning based on new data, incident experience and interaction with the adversary. Automation of routine risk assessment tasks, prioritisation, and even certain elements of response (with human oversight of critical decisions) will become indispensable to maintaining effectiveness in the face of the scale and pace of attacks.

Despite methodological advances, managing under conditions of deep uncertainty will remain a constant challenge. The fundamental unpredictability of the emergence of entirely new types of threats means that stochastic cyber risk management models must be resilient to this uncertainty. This will require combining quantitative methods with qualitative scenario analysis, applying robust optimisation techniques, and making extensive use of Bayesian approaches that actively update beliefs about the state of risk as new information becomes available. The concept of cyber resilience – understood as an organisation's ability to continue to fulfil its mission and basic functions despite a serious cyber incident – comes to the fore as an overarching goal to which stochastic cyber risk management methods should contribute.

Ethics and responsibility must be the foundation for the development and implementation of advanced stochastic cyber risk management technologies. The principles of transparency, fairness, respect for privacy and meaningful human oversight must be built into the design of systems from the outset. The development of a legal and regulatory framework that addresses the issue of responsibility for decisions made by autonomous systems and ensures adequate protection of individual rights is urgently needed.

The conclusion is clear: stochastic cyber risk management is no longer an academic curiosity, but has become a necessary foundation for building truly resilient organisations in an increasingly dangerous and unpredictable cyberspace. Despite significant challenges related to data, computation, interpretability, model security, and ethical and legal aspects, the directions for development outlined in this paper offer a powerful and increasingly mature arsenal of tools for quantitatively understanding, predicting, and actively shaping an organisation's risk profile. Investment in further research, methodological development and, most importantly, the practical implementation of these methods in organisations across various sectors is crucial not only for the security of individual entities, but also for national security, economic stability and the protection of fundamental rights, such as privacy, in the coming decades. The future belongs to a quantitative, adaptive and interdisciplinary approach to cyber risk.

Bibliography

- **Alfredsson, P., & Verendel, V. (2017).** Cyber Security–Stochastic Models and Applications. W: *Stochastic Models in Reliability Engineering* (s. 317-340). CRC Press. (Podstawy modelowania stochastycznego w cyberbezpieczeństwie).
- **Eling, M., & Wirfs, J. H. (2019).** What are the actual costs of cyber risk events? *European Journal of Operational Research*, 272(3), 1109-1119. (Kwantyfikacja finansowych skutków zdarzeń cybernetycznych).
- **Woods, D. W., & Moore, T. (2020).** Does Insurance Have a Future in Governing Cybersecurity? *IEEE Security & Privacy*, 18(1), 21-27. (Krytyczna analiza roli ubezpieczeń).
- **Apruzzese, G., et al. (2023).** The Role of Machine Learning in Cybersecurity. *Digital Threats: Research and Practice*, 4(1). (Przeglądowy artykuł dot. zastosowań ML).
- **Bridges, R. A., et al. (2019).** Probabilistic Learning of the Exact Topology of a Bayesian Network. *Proceedings of the IEEE International Conference on Big Data*. (O probabilistycznych sieciach Bayesowskich).
- **Goodfellow, I., Bengio, Y., & Courville, A. (2016).** *Deep Learning*. MIT Press. (Podręcznik fundamentów głębokiego uczenia, w tym GAN i VAE).
- **Hindy, H., et al. (2020).** A Taxonomy of Network Threat Hunting Approaches. *IEEE Access*, 8, 130359-130375. (Taksonomia obejmująca modele anomalii).
- **Gilbert, N., & Troitzsch, K. G. (2005).** *Simulation for the Social Scientist*. Open University Press. (Podstawy modelowania agentowego).

- **Grimm, V., et al. (2020).** The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism. *Journal of Artificial Societies and Social Simulation*, 23(2). (Standard opisu modeli ABM).
- **Nguyen, T. H., & Reddi, V. J. (2021).** Deep Reinforcement Learning for Cyber Security. *IEEE Transactions on Neural Networks and Learning Systems*. (Połączenie RL i cyberbezpieczeństwa).
- **Sawilla, R. E., & Wiemer, D. (2018).** A System-Aware Cyber Resilience Model for Mission Assurance. *Proceedings of the IEEE Military Communications Conference (MILCOM)*. (Zastosowanie do analizy odporności).
- **Adkins, G., et al. (2021).** The DARPA GARD Program. arXiv preprint arXiv:2111.00566. (Program badawczy dot. odporności modeli ML na ataki).
- **Barocas, S., Hardt, M., & Narayanan, A. (2019).** Fairness and Machine Learning: Limitations and Opportunities. *fairmlbook.org*. (Podręcznik dot. uczciwości i biasu w ML).
- **European Union Agency for Cybersecurity (ENISA). (2021).** AI Cybersecurity Challenges: Threat Landscape for Artificial Intelligence. (Raport dot. wyzwań bezpieczeństwa AI).
- **Goodman, B., & Flaxman, S. (2017).** European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation". *AI Magazine*, 38(3). (Podstawy prawne dla XAI w UE).
- **Veale, M., Van Kleek, M., & Binns, R. (2018).** Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. (Etyka i odpowiedzialność w systemach autonomicznych).
- **Coles, S. (2001).** An Introduction to Statistical Modeling of Extreme Values. *Springer Series in Statistics*. (Podstawowy podręcznik EVT).
- **Hurd, T. R. (2016).** Contagion! Systemic Risk in Financial Networks. Springer. (Modelowanie ryzyka kaskadowego w sieciach – koncepcje applicable do cyber).
- **World Economic Forum (WEF). (2023).** The Global Risks Report. (Coroczny raport identyfikujący ryzyka systemowe, w tym cyber).
- National Institute of Standards and Technology (NIST). (2018). Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (Ramowy dokument dla zarządzania ryzykiem, ewoluujący w kierunku integracji z metodami ilościowymi).
- **Zuech, R., Khoshgoftaar, T. M., & Wald, R. (2015).** Intrusion detection and big heterogeneous data: a survey. *Journal of Big Data*, 2(1). (Wyzwania związane z danymi w cyber).
- **The EU Agency for Cybersecurity (ENISA). (2023).** *Horizon 2025: Emerging Cyber-Threats and Threat Intelligence*. (Prognoza przyszłych trendów i zagrożeń).