Using Economic Theory to Dissect Cyber Espionage

Laiba Khalid

ABSTRACT

This paper applies economic theory to analyze market failures in AI-driven cybersecurity, focusing on the illicit market for stolen data. It examines foundational concepts such as externalities, information asymmetry, transaction costs, and behavioral economics to explain structural drivers of cyber insecurity. Using the 2021 Microsoft Exchange Server hack as a case study, the paper demonstrates how these economic forces manifest in real-world cyberattacks, including the disproportionate impacts on firms, customers, and society. The analysis highlights how misaligned incentives, opaque markets, and concentrated technological power exacerbate vulnerability, here understood as the economically conditioned susceptibility of systems and actors to attack, and reduce the effectiveness of defense measures. Building on this framework, the paper identifies policy interventions aimed at realigning incentives, improving transparency, and fostering collaborative defense strategies in digital ecosystems. By integrating economic reasoning with cybersecurity analysis, it provides a systematic approach for understanding and addressing persistent market failures in the increasingly AI-mediated digital environment.

https://doi.org/10.4079/pp.v33i0.11

Laiba Khalid is a second-year MPP candidate with a concentration in Science and Technology Policy. Her research focuses on AI governance and emerging policy challenges. She currently works as a Communications Assistant at the Trachtenberg School of Public Policy and Public Administration. Originally from Pakistan, she brings two years of experience in the media and policy sectors. In her free time, she enjoys cooking and watching movies.

ACKNOWLEDGEMENTS

The author expresses her sincere gratitude to Professors Anil Nathan, Carol Kuntz, and Kathy Newcomer for their guidance and thoughtful feedback on earlier drafts of this paper. She would also like to extend her appreciation to her colleagues, Christian Schamberger and Thakur Bhanu Partap, for their helpful edits and valuable recommendations. Lastly, she acknowledges that any remaining errors are her own.

Introduction

Artificial intelligence (AI) has emerged as both sentinel and saboteur in the digital arms race of the 21st century. It empowers cybersecurity teams to detect or even predict anomalies and breaches in real-time, and automates defenses at a scale that was once unimaginable. But, in the hands of adversaries it becomes a weapon of precision, supercharging phishing schemes, automating network infiltration, and scaling cyberattacks beyond human limitations. In 2023 alone, breach success rates for AI-enhanced phishing campaigns increased manifold, fueling an underground data economy now valued at more than a billion dollars (Europol 2023). Cybersecurity Ventures predicts the global cybercrime cost to reach $10.5 trillion by 2025. Thus, the horizon for cyber espionage has expanded from a state-centered covert operation to a decentralized, highly profitable enterprise accessible to both nation-states and private actors.

However, beneath the technological arms race lies a deeper problem: the economic architecture of cybersecurity is fundamentally broken. Firms routinely underinvest in protection not because they lack tools, but because cybersecurity generates both positive and negative externalities. Investments in security create spillover benefits for interconnected actors, while breaches impose costs on customers, partners, and society that firms do not fully internalize, leading to systematic underinvestment (Kunreuther and Heal 2003). Meanwhile, attackers exploit zero-day vulnerabilities, flaws that remain unknown and unpatched at the time of attack, and often possess crucial information long before anyone else becomes aware of it, leading to systemic asymmetry in risk exposure (Akerlof 1970). Outsourcing security introduces principal-agent conflict, as vendors prioritize efficiency and cost-cutting over resilience (Williamson 1981). Even when organizations are aware of threats, psychological biases such as optimism bias, where risks are systematically underestimated, present bias, which leads decision-makers to prioritize short-term costs over long-term security investments, and the tendency to discount low-probability but high-impact events result in dangerous delays in action (Kahneman and Tversky 1979).

This paper explains how the growing threat of AI-driven cyber espionage cannot be understood through technology alone. It could more effectively be viewed through the lens of microeconomic theory, which reveals how misaligned incentives, information asymmetry, cognitive biases, and market imperfections shape security outcomes. Analysis of the 2021 Microsoft Exchange Server breach illustrates how core economic failures converge to amplify risk and delay response. Finally, the paper concludes with targeted policy interventions that seek to realign market incentives and harness AI's potential without stifling

innovation.

The Economic Distortions in AI-Driven Cybersecurity

Externalities and the Underinvestment Problem

Underinvestment in cybersecurity is one of the most persistent economic failures in the field, driven by firms' inability to fully recognize or capture their value. While cybersecurity investments generate significant private benefits, they also produce broader spillover effects for networks, supply chains, and national infrastructure by reducing systemic risk. However, because firms cannot internalize these wider benefits, they have weaker incentives to invest at socially optimal levels. This results in chronic underinvestment, leaving interconnected systems vulnerable to cascading failures. A striking example of the catastrophic costs that can arise from such underinvestment is the 2017 NotPetya cyberattack.

NotPetya originated as a targeted disruption campaign against Ukrainian infrastructure before quickly spiraling into a global catastrophe. Disguised as ransomware and propagated via a popular accounting software update, NotPetya exploited a vulnerability in Microsoft systems called "EternalBlue" (Greenberg 2019). The malware rendered infected systems entirely inoperable by overwriting critical data, making recovery impossible without full system rebuilds, destroying operational capacity rather than merely stealing information. Within hours, NotPetya infected multinational firms far from the conflict zone. Shipping giant A.P. Møller-Mærsk was forced to halt port operations worldwide, as the malware wiped the IT systems underpinning its entire global logistics network. Pharmaceutical company Merck & Co. lost access to critical manufacturing systems. FedEx's European subsidiary TNT Express experienced severe logistical breakdowns. A White House assessment estimated the total economic damage from NotPetya at more than $10 billion, making it one of the most destructive cyberattacks in history (Steinberg et al. 2021).

The NotPetya incident highlights the consequences of ignoring positive externalities. The software vendor rationally underinvested in secure design based on private costs and benefits, but failed to account for the massive social benefits a secure product would have generated. In economic terms, the marginal private benefit (MPB) of investing in defense is substantially lower than the marginal social benefit (MSB), which includes the cost of preventing spillover damage. Consequently, the level of investment in security chosen by

private actors is consistently lower than the socially optimal level.

This misalignment between private and social incentives generates a collective action problem: since each firm benefits from the cybersecurity investments of others but has limited incentive to contribute proportionally, the result is chronic under-protection at the system level (Anderson and Moore 2006; CNIL 2025).

In other words, the benefits of strong cybersecurity extend far beyond the boundaries of the investing firm. And without proper compensation, companies continue to rationally underinvest. The social cost of this under-provision is only set to rise as AI enables increasingly fast, precise, and scalable cyberattacks.

Information Asymmetry

Efficient markets rely on transparency. When buyers and sellers have equal access to information, prices accurately reflect quality, and rational choices prevail. However, in the domain of cybersecurity, the stark imbalance of knowledge consistently skews incentives and outcomes. IBM's 2024 X-Force Threat Intelligence Index finds that attackers frequently rely on stolen credentials and the exploitation of public-facing applications, often using known vulnerabilities with publicly available exploit code, rather than novel zero-days (previously unknown software flaws for which no patch yet exists). This dynamic creates structural advantage: defenders must secure entire systems continuously, while attackers need to only identify and exploit a single overlooked weakness (IBM 2024).

AI has deepened the asymmetry between attackers and defenders. Tools that can autonomously scan for vulnerabilities and tailor attacks to specific system configurations make offensive capabilities easier and more manageable. Research in adversarial machine learning shows that attackers can use automated techniques, such as reinforcement learning-based adversarial packet generators, to systematically bypass model-based defenses, illustrating how AI tools themselves can be weaponized against defenders (Hore et al. 2023).

George Akerlof's concept "Market for Lemons" (1970) explains the economic implications of such information asymmetries. In markets where sellers possess more information than buyers, high-quality goods tend to disappear because buyers are unwilling to pay premiums for products whose value they cannot

verify. This dynamic is visible in consumer markets such as used cars or health insurance, where buyers routinely struggle to distinguish genuine quality from surface-level presentation. The result is adverse selection, where lower-quality products dominate and trust deteriorates.

In cybersecurity, the parallel is striking. Security service consumers (especially small firms) often cannot assess on their own the quality or comprehensiveness of protection offered by vendors. This leads to underinvestment in better security tools and an overreliance on brand names or surface-level assurances. At the same time, firms can also develop a false sense of security, overestimating the resilience of their systems simply because they have not yet been breached. In most cases risks are invisible until it's too late.

This asymmetry in information has an impact on the entire market. High quality cybersecurity providers are unable to advertise their competencies due to the absence of standardized metrics or third-party evaluations. The attackers, on the other hand, only need to get in once to impose their authority. Moreover, the opacity of proprietary software, like Microsoft Exchange, reduces the ability of third-party researchers to audit or detect hidden vulnerabilities, further distorting the balance (Kshetri 2010).

Some policymakers have advocated mandatory vulnerability disclosure laws and increased transparency requirements. For example, the European Union's Cyber Resilience Act mandates firms to report significant vulnerabilities to national cybersecurity authorities within 24 hours of discovery (European Commission 2024). Such policies could help align information access and encourage preemptive defense strategies across the market.

Transaction Costs and the Principal-Agent Problem

As threats rise, many organizations outsource cybersecurity to specialized AI-enabled security firms, vendors that deploy machine learning, automated threat detection, and predictive analytics to manage defense at scale (IBM 2024). This offers technical advantages but also exposes new risks. The core issue being the principal-agent problem. The principal (government or corporation) wants maximum protection. The agent (cybersecurity firm) may prioritize profits. When incentives diverge, security suffers.

This misalignment in incentives is amplified by transaction costs. Contracts are often complex and difficult to enforce. Monitoring security providers is

costly and sometimes infeasible. Compliance costs alone consume up to 30% of cybersecurity budgets (Anderson, 2020).

These structural inefficiencies become visible in real-world cybersecurity partnerships and technology procurement decisions. For example, observers have raised concerns that Palantir's proprietary platforms, by virtue of vendor lock-in, a condition in which clients become so dependent on a single provider's architecture that switching to alternative systems becomes prohibitively costly, technically complex, or both. This dependency, combined with limited interoperability with external systems, can inhibit data portability and collaborative information sharing, potentially creating functional silos that hinder broader analytical cooperation (Hash.ai 2025). In situations like these, the true cost of outsourcing isn't just financial. It includes weakened national or corporate security.

Oliver Williamson's theory of transaction cost economics (1981) offers relevant propositions. For example, when transaction costs are high, markets may fail to allocate resources efficiently. In such conditions, vertical integration is a way forward where the government builds in-house AI defense systems that could be more efficient than private solutions.

Behavioral Economics

Behavioral economics offers crucial insight into why firms systematically underprepare for cyber threats. A firm's decision-making process plays a great role in cybersecurity. Even in cases of perfect information, firms often fail to act decisively and delays in acting highlight the deep-rooted cognitive biases of decision-makers.

Prospect theory says that individuals tend to weigh losses more heavily than equivalent gains (Kahneman and Tversky 1979). In cybersecurity, loss aversion manifests in two ways. First, firms are reluctant to invest large sums in preventive security measures because the benefits are invisible. i.e., measured in events that "do not" happen. Second, the perceived loss of capital today outweighs the potential gain of avoiding an uncertain future breach.

Decision-making flaws are further fueled by optimism bias and the availability heuristic, which is defined as the cognitive tendency to assess the likelihood of an event based on how easily a similar example comes to mind, rather than on objective probability. Executives presume that their organizations are less

likely to be targeted as compared to their competition. Additionally, a breach appears real only if it is recent or high-profile. According to a Europol study, many CEOs delayed major AI-security investments until their firms suffered a cyberattack, a pattern consistent with reactive rather than proactive risk management (Europol 2024).

Cognitive biases affecting decisions to invest in cybersecurity can have significant consequences. Cybersecurity becomes a "regret-driven" domain, in which investments are more likely to occur after a loss, instead of being made for its prevention. Such decisions are economically inefficient and socially costly, especially given the networked nature of cyber risk, where one firm's delay can endanger an entire supply chain.

Market Structure and Oligopoly Power in AI Cybersecurity

The market structure of AI-driven cybersecurity is characterized by significant oligopoly power, contributing to persistent inefficiencies in both innovation and investment. It is dominated by a small number of technology firms, namely Google, Microsoft, and Amazon. These firms control a disproportionate share of intellectual property related to AI threat detection. The Federal Trade Commission has raised significant concerns about market concentration and the potential for anti-competitive practices in the AI sector, which can stifle innovation in areas like cybersecurity, in addition to making it difficult for smaller rivals to enter the market (FTC 2023).

This market dominance is reinforced by several structural elements. Firstly, smaller businesses find it challenging to compete due to the substantial economies of scale created by the high fixed costs of research and development in AI cybersecurity, particularly for training large language models and prediction systems (Anderson and Moore 2006). Second, for AI-based threat detection to be effective, access to vast amounts of real-time data is essential. Dominant companies frequently treat these data as confidential, limiting new entrants' access (Varian 2004). Furthermore, "patent thickets" stifle the possibility of disruptive innovation by legally solidifying incumbents' positions (Brangetto and Kert-Saint Aubyn 2015).

These dynamics foster regulatory capture where big AI companies have the financial and political clout to shape cybersecurity laws, frequently pushing for guidelines that complement their own patent portfolios and technological architecture (OECD 2022). The dominance of a few firms discourages innovation at grassroots level and makes regulatory agencies less responsive

to the demands of the larger market. In contrast to industries with more fragmented participation, the concentration of AI development among a few dominant firms may slow the broader adoption of robust security practices. While many organizations recognize cybersecurity as a key AI risk, as the Stanford HAI 2025 AI Index Report notes, a substantial proportion still do not actively mitigate it, reflecting an innovation and implementation gap rather than effective competitive discipline.

Case Study: Microsoft Exchange Server Hack (2021)

The 2021 Microsoft Exchange Server hack marked one of the most widespread and damaging cyberattacks in recent history, compromising the email systems of over 30,000 U.S. organizations including government agencies, small businesses, and major corporations (Krebs 2021). The attackers, identified as Hafnium, exploited four zero days in Microsoft's Exchange software to gain unauthorized access to email accounts, exfiltrate data, and plant backdoors for continued access even after detection. This breach was particularly devastating due to the speed with which it spread before the organization could even respond. Microsoft eventually released patches, but by that point, the damage had been done; thousands of servers had been compromised, and a costly remediation process ensued. This incident reveals how structural economic failures in cybersecurity exacerbate both the frequency and impact of cyberattacks. By analyzing the breach through an economic lens, it becomes clear that technological vulnerabilities are not the sole drivers of modern cyber crises; rather, they are symptoms of deeper market inefficiencies and flawed incentives that allow such attacks to flourish (Anderson and Moore 2006; Akerlof 1970; Williamson 1981; Kahneman and Tversky 1979).

Negative externalities can be seen at the heart of this breach. While Microsoft's losses were primarily reputational and limited to the costs of patch development, their clients bore the full breadth of the breach's consequences; financial, operational, and institutional. Estimated client losses reached $2.5 billion (Newman 2021), but the damage extended well beyond balance sheets. Many public institutions experienced email disruptions and temporary service outages that affected the communities they served, particularly in settings where institutions cannot afford in-house security expertise. Hospitals faced interruptions to patient communication systems, municipalities lost access to administrative infrastructure, and small businesses suffered data losses with no clear path to recovery (Krebs 2021; FBI and CISA 2021). Microsoft's responsibility was limited by current legal and regulatory frameworks, allowing it to underinvest in proactive vulnerability testing and secure design practices. The costs borne by municipalities, hospitals, and small businesses, i.e. entities

entirely dependent on Microsoft software with no control over its underlying vulnerabilities, illustrate the large-scale human and economic impacts of these externalities on clients and bystanders.

Looking at information asymmetry from the perspective of this exchange hack, especially the kind outlined in Akerlof's model, markets deteriorate when one party consistently possesses more or better information than the other. As already discussed, in cybersecurity, the imbalance between attackers and defenders is profound. Hafnium exploited zero-day vulnerabilities, the flaws that were unknown to Microsoft and its users at the time of the attack. Such vulnerabilities often remain undiscovered for extended periods; in fact, the average zero-day exploit is said to have remained active for over 300 days (Forbes 2012). During that time, attackers operated with impunity while defenders remained in the dark, unaware that their systems were compromised.

Microsoft Exchange Software's proprietary nature also fuels its hidden vulnerabilities. Open-source platforms allow third-party researchers to identify threats or flaws. However, Microsoft operates on closed-source systems, which restrict scrutiny to internal teams. This limitation fosters a lack of transparency, where software vendors maintain informational dominance but bear insufficient accountability for the consequences of unknown flaws. Even when Microsoft released patches after discovering the breach, many organizations were unable to act properly due to a lack of technical expertise or resources. This delay in response widened the attack's reach and prolonged the damage.

Further compounding the Microsoft Exchange Server Hack problems was the principal-agent dilemma embedded in how organizations manage IT security. Many organizations affected by the Exchange Server breach outsourced their IT management to third-party vendors or managed service providers (MSPs). These vendors were responsible for patching systems, but their incentives were not always aligned with those of their clients. For example, some MSPs delayed patching due to operational constraints and compatibility testing requirements (Rapid7 2021). Others had limited capacity to respond to a large-scale vulnerability across multiple clients simultaneously.

This separation of responsibility (agent) from risk-bearing (principal) created a situation in which critical updates were not prioritized appropriately. Clients, particularly small organizations and local governments, were often unaware of the urgency or unable to enforce timely patching. This gap shows the inefficiencies introduced by transaction costs and contracting challenges in cybersecurity services. The process of coordination with external service providers is complex, and the cost of ensuring consistent patch management is

high even for well-intentioned organizations.

An additional insight from the Exchange Server case is the role of behavioral biases in shaping organizational responses to cybersecurity threats. Traditional economic models assume rational actors, but decades of research suggest that decision-making is influenced by cognitive biases and heuristic shortcuts. In this case, many organizations displayed the "it won't happen to us" bias, underestimating their exposure to high impact cyberattacks. This misperception can result in the delay of critical updates, and operation under the false assumption that state-sponsored attacks target only large enterprises or government institutions.

Furthermore, the status quo bias, which is an aversion to change even when the current state is risky, also played a role. IT teams, wary of the disruptions and compatibility issues that patches might cause, often postponed updates in favor of short-term operational stability. This procrastination proved catastrophic, as the window between patch release and widespread exploitation was very narrow. In this regard, the breach was less a failure of knowledge than of action, driven by the very human tendency to discount abstract future risks in favor of present convenience.

Lastly, the Microsoft Exchange breach illustrates the danger of market concentration in software ecosystems. The fact that Microsoft had such a dominant position in enterprise email solutions says a lot about the high and far-reaching implications this breach had. Unlike competitive markets, where the failure of one provider might affect only a subset of users, Microsoft's market share amplified the systemic risk of a single point of failure. This monoculture creates fragility, as attackers rapidly exploit interconnected systems with a single toolset. From an economic standpoint, the lack of competitive pressure reduced Microsoft's incentives to ensure a resilient security infrastructure. The case thus provides a cautionary tale about how monopolistic software ecosystems can amplify cybersecurity risks across sectors and borders.

Policy Recommendations

Based on the distortions discussed in the paper, six economic policy solutions are proposed:

1. Taxation

A Pigouvian tax would require firms to pay a levy proportional to their contribution to systemic cyber risk, assessed prospectively rather than retroactively. Critically, this is not a fine levied after a breach occurs. Rather, it is an ex ante regulatory instrument calibrated to a firm's ongoing risk profile: the size of its attack surface, the sensitivity of the data it holds, its degree of interconnectedness with critical infrastructure, and the adequacy of its existing security controls. Established cyber risk quantification frameworks such as FAIR (Factor Analysis of Information Risk) and the NIST SP 800-30 risk assessment standard already provide methodologies for translating these variables into monetized risk estimates, offering regulators a workable basis for setting tax rates (NIST 2012). Analogous mechanisms have been proposed in financial markets, where Pigouvian-style systemic risk taxes on financial institutions are calibrated to each firm's contribution to cascade risk within interconnected networks (Zlatić et al. 2015). The revenue generated could fund public threat intelligence infrastructure or subsidize security upgrades for under-resourced organizations. This mechanism would internalize the social costs of cybersecurity failures and incentivize firms to invest at levels closer to the socially optimal equilibrium.

2. AI Security Tools Subsidy:

Many firms, especially small and mid-sized enterprises, underinvest in cybersecurity because they face resource constraints and anticipate that others will provide the protection from which they can indirectly benefit. Public subsidies for open-access AI-driven security tools such as anomaly-detection algorithms or secure-by-design code libraries would reduce these free-rider problems by creating shared, high-quality defenses accessible to all market participants. By lowering the marginal cost of adoption, such subsidies enable more uniform baseline security across the economy and prevent disproportionate vulnerabilities among smaller firms.

3. Public Threat Intelligence Hubs:

Government-led threat intelligence platforms, modeled on the Cybersecurity and Infrastructure Security Agency's (CISA) 2023 initiatives, facilitate real-time exchange of indicators of compromise, attacker tactics, and AI-driven threat analytics (Duffy 2023). These hubs correct information asymmetries by reducing the gaps between well-resourced firms and those with limited monitoring capacity. They also strengthen collective defense by distributing knowledge of emerging attack vectors more broadly and rapidly across sectors.

Behavioral Nudges: Behavioral interventions can systematically counteract the cognitive biases that lead firms to underprepare for cyber threats. Cyber risk calculators present organizations with quantified, scenario-specific estimates of their financial exposure, translating abstract probabilities into concrete dollar-value losses that are harder to discount than general warnings. Gamified training programs embed security awareness into employees' daily workflows through simulations and point-based rewards. Research has found that gamified security training significantly improves employees' security self-efficacy and measurably reduces rates of falling victim to phishing simulations (Bitrián et al. 2024). Post-breach simulations expose executives to realistic reconstructions of attack scenarios, activating the availability heuristic constructively by making the threat feel proximate and concrete rather than distant. These tools would be directed at firms and their decision-makers. Mandated disclosure of standardized risk scores to regulators and in some cases the public may further incentivize precautionary action by creating reputational and financial consequences for persistent underinvestment.

4. Antitrust Enforcement in AI Security:

AI cybersecurity markets increasingly exhibit concentration, where a few dominant firms provide vertically integrated systems that are difficult to replace or interoperate with. Such vendor lock-in effects arise when switching costs are high, because data formats, models, or security architectures are proprietary, thereby limiting competition and innovation in defensive technologies. The Federal Trade Commission should scrutinize exclusionary practices, promote open standards, and encourage modular security architectures i.e. designs in which individual security components such as threat detection and response systems are built as interchangeable, independently upgradeable units that communicate through standardized interfaces rather than being fused into a single proprietary stack (NIST 2018). Reducing lock-in would lower barriers to entry for new firms, increase market competition, and accelerate improvements in AI-enabled defenses.

5. Cybersecurity Rating Systems for Vendors:

Like credit scores for finance, cybersecurity ratings would reduce transaction costs by giving clients a reliable, low-cost signal of vendor quality. This directly addresses the Akerlof-style information asymmetry identified earlier in the paper. A robust ratings regime, administered by an independent body and based on auditable criteria derived from frameworks such as NIST CSF or ISO 27001, would increase market transparency and create competitive pressure on vendors to invest in genuine resilience rather than surface-level assurances (NIST 2018).

Conclusion

To sum up, this paper demonstrated how microeconomic theory can be used to analyze failures in the cybersecurity market, especially with the advent of AI. The analysis shows that asymmetric information enables threat actors to exploit system weaknesses undetected, misaligned private and social incentives allow organizations to dump security costs onto society, and high transaction costs discourage proactive defense investments. Additionally, behavioral biases influence decision-makers to underinvest in long-term cyber resilience. The Microsoft Exchange Server Hack case study illustrated how these dynamics play out in wide-reaching real-world breaches, affirming the need for stronger regulatory intervention and incentive alignment.

Future research should focus on developing empirical tools to quantify the economic costs of AI-driven cyberattacks and evaluate the effectiveness of different policy responses. Building standardized datasets that capture both direct damages and secondary effects, such as reputational harm or supply chain disruption, will be essential. Moreover, as AI continues to evolve, interdisciplinary approaches, combining economics, computer science, evaluation, and behavioral psychology will be critical to understanding and mitigating its risks. Implementing the recommendations offered here could help transition cybersecurity from a reactive posture to a more predictive and economically efficient framework.

References

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." The Quarterly Journal of Economics 84 (3): 488–500. https://doi.org/10.2307/1879431

Anderson, Ross. 2020. Security Engineering: A Guide to Building Dependable Distributed Systems. 3rd ed. Hoboken, NJ: Wiley.

Anderson, Ross, and Tyler Moore. 2006. "The Economics of Information Security." Science 314 (5799): 610–13. https://doi.org/10.1126/science.1130992

Bitrián, Paula, Isabel Buil, Sara Catalán, and Dominik Merli. "Gamification in Workforce Training: Improving Employees' Self-Efficacy and Information Security and Data Protection Behaviours." Journal of Business Research 179 (2024): 114685. https://doi.org/10.1016/j.jbusres.2024.114685

Brangetto, Paolo, and Michel Kert-Saint Aubyn. 2015. Economic Aspects of National Cyber Security Strategies. Project Report. NATO Cooperative Cyber Defence Centre of Excellence. https://afyonluoglu.org/PublicWebFiles/library/ccdcoe/LIB_0021.pdf

CISA (Cybersecurity and Infrastructure Security Agency). Remediations for Microsoft Exchange Server Vulnerabilities. Washington, DC: CISA, 2021. https://www.cisa.gov/news-events/alerts/2021/03/08/cisa-mitigations-and-workarounds-published-for-microsoft-exchange

Commission Nationale de l'Informatique et des Libertés (CNIL). 2025. "Cybersecurity: The Economic Benefits and Externalities." https://www.cnil.fr/en/cybersecurity-economic-benefits-gdpr

Duffy, Michael. 2023. "Enabling Threat-Informed Cybersecurity: Evolving CISA's Approach to Cyber Threat Information Sharing." Cybersecurity and Infrastructure Security Agency, December 18, 2023. https://www.cisa.gov/news-events/news/enabling-threat-informed-cybersecurity-evolving-cisas-approach-cyber-threat-information-sharing

Europol. 2024. Internet Organised Crime Threat Assessment (IOCTA) 2024. https://www.europol.europa.eu/cms/sites/default/files/documents/Internet%20Organised%20Crime%20Threat%20Assessment%20IOCTA%202024.pdf

European Commission. "Cyber Resilience Act: Regulation on Cybersecurity Requirements for Products with Digital Elements." European Commission, 2024. https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act

FBI and CISA. Compromise of Microsoft Exchange Server. IC3 Joint Cybersecurity Advisory, 2021. https://www.ic3.gov/CSA/2021/210310.pdf

Federal Trade Commission (FTC). 2023, June 29. "Generative AI Raises Competition Concerns." https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns

Greenberg, Andy. 2012. "Hackers Exploit 'Zero-Day' Bugs for 10 Months on Average Before They're Fixed." Forbes, October 16, 2012. https://www.forbes.com/sites/andygreenberg/2012/10/16/hackers-exploit-software-bugs-for-10-months-on-average-before-theyre-fixed/

Greenberg, Andy. 2019. Sandworm: A New Era of Cyberwar and the Hunt for the Kremlin's Most Dangerous Hackers. https://kclpure.kcl.ac.uk/ws/portalfiles/portal/257062100/Sandworm-_A_New_Era_of_Cyberwar_and_the_Hunt_for_the_Kremlin_s_Most_Dangerous_Hackers.pdf

Hash.ai. "The Problem with Palantir." Hash.ai Blog, 2025. https://hash.ai/blog/the-problem-with-palantir

Heal, Geoffrey, and Howard Kunreuther. 2004. Interdependent Security: A General Model. NBER Working Paper 10706. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w10706

Hore, Soumyadeep, Jalal Ghadermazi, Diwas Paudel, Ankit Shah, Tapas K. Das, and Nathaniel D. Bastian. 2023. "Deep PackGen: A Deep Reinforcement Learning Framework for Adversarial Network Packet Generation." arXiv (May 18, 2023). https://doi.org/10.48550/arXiv.2305.11039

IBM. 2024. IBM X-Force Threat Intelligence Index 2024. Armonk, NY: IBM Corporation. https://www.ibm.com/think/x-force/2024-x-force-threat-intelligence-index

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." Econometrica 47 (2): 263–91. https://doi.org/10.2307/1914185

Kshetri, Nir. 2010. "An Institutional Perspective on Cybercrimes." In The Global Cybercrime Industry, 57–74. Springer. https://ideas.repec.org/h/spr/sprchp/978-3-642-11522-6_3.html

Krebs, Brian. 2021, March 5. "At Least 30,000 U.S. Organizations Newly Hacked via Holes in Microsoft's Email Software." Krebs on Security. https://krebsonsecurity.com/2021/03/at-least-30000-u-s-organizations-newly-hacked-via-holes-in-microsofts-email-software/

Newman, Lily Hay. 2021, March 8. "The Microsoft Exchange Hack Is Worse Than First Thought." Wired. https://www.wired.com/story/microsoft-exchange-patch-hacks-ransomware/

NIST (National Institute of Standards and Technology). Guide for Conducting Risk Assessments: SP 800-30 Rev. 1. Gaithersburg: NIST, 2012. https://doi.org/10.6028/NIST.SP.800-30r1

NIST (National Institute of Standards and Technology). Framework for Improving Critical Infrastructure Cybersecurity (CSF) Version 1.1. Gaithersburg: NIST, 2018. https://doi.org/10.6028/NIST.CSWP.04162018

OECD. 2022. Digital Security Policy Framework. Paris: OECD. https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/12/oecd-policy-framework-on-digital-security_a0b1d79c/a69df866-en.pdf

Rapid7. (2021). "For Microsoft Exchange Server Vulnerabilities, Patching Remains Patchy." Rapid7 Blog, October 6. https://www.rapid7.com/blog/post/2021/10/06/for-microsoft-exchange-server-vulnerabilities-patching-remains-patchy/

Stanford Institute for Human-Centered AI (HAI). 2025. AI Index Report 2024. Stanford University. https://hai-production.s3.amazonaws.com/files/hai_ai_index_report_2025.pdf

Steinberg, Sean, Adam Stepan, and Kyle Neary. 2021. NotPetya: A Columbia University Case Study. New York: Columbia University School of International and Public Affairs, Picker Center Digital Education Group. https://www.sipa.columbia.edu/sites/default/files/2022-11/NotPetya%20Final.pdf

Varian, Hal R. 2004. "System Reliability and Free Riding." In Economics of Information Security, edited by L. Jean Camp and Stephen Lewis, 1–15. Advances in Information Security 12. Boston, MA: Springer. https://link.springer.com/chapter/10.1007/1-4020-8090-5_1

Williamson, Oliver E. 1981. "The Economics of Organization: The Transaction Cost Approach." American Journal of Sociology 87 (3): 548–77. https://doi.org/10.1086/227496

Zlatić, Vinko, Giampaolo Gabbi, and Hrvoje Abraham. "Reduction of Systemic Risk by Means of Pigouvian Taxation." PLOS ONE 10, no. 7 (2015): e0114928. https://doi.org/10.1371/journal.pone.0114928