Cybersecurity at MIT Sloan

Cybersecurity at MIT Sloan Annual Members Only Conference

Cybersecurity and the Next Wave of AI Systems

June 3, 2026 | In-Person & Virtual via Zoom

Event Overview

Cybersecurity at MIT Sloan (CAMS) expert researchers examine autonomous systems, digital risk, and emerging security challenges.

Artificial intelligence is no longer just generating outputs, it is beginning to act.
Across enterprise systems, AI is moving into roles that initiate decisions, execute workflows, and influence real-world outcomes with limited human oversight. As this shift accelerates, long-standing cybersecurity assumptions are starting to break.

Systems are no longer static. Trust boundaries are no longer clear. And in many cases, organizations are deploying AI faster than they can secure it.
This year's annual conference brings together researchers, industry leaders, and policymakers to confront a central question: What does it mean to secure systems that can think, generate, and act on their own?

Participants will engage with questions that are rapidly moving from theoretical to operational:

- What happens when AI systems act on unverified or manipulated inputs?
- Where do new attack surfaces emerge, and existing surfaces reshape when AI systems generate and execute their own outputs?
- How do you secure a supply chain when the intelligence inside it can be influenced or poisoned?
- What governance models are viable when decision-making is partially delegated to machines?

This is not a future problem. These dynamics are already unfolding across enterprise environments. The organizations that understand this shift early will be better positioned to manage risk, build resilience, and shape how AI systems are deployed securely at scale

Program Overview

9:00AM — Conference Overview & Goals: Opening Remarks & Welcome (15 min Opening Remarks) Professor Stuart Madnick & Dr. Michael Siegel (MIT CAMS Directors)

9:15AM — The Escalating Agentic AI Arms Race between the Defenders and Attackers (45 min Academic Talk) Professor Stuart Madnick, MIT CAMS Founder

10:00AM — **Offensive AI Is Here. What Do We Do About It?** (45 min Academic Keynote) Dr. Yisroel Mirsky, Zuckerman Faculty Scholar, Head of the Offensive AI Lab, Institute of Software, Systems and Security, Ben-Gurion University

10:45AM — Morning Break (30 min)

11:15AM — When Cyber Risk Becomes Non-linear, Governance Must Become Adaptive (30 min Academic Talk) Dr. Sander Zeijlemaker, Research Affiliate, MIT CAMS

11:45AM — **Poisoned Intelligence: How LLMs Are Reshaping Software Supply Chain Cybersecurity** (60 min Panel)
Moderator: Dr. Ranjan Pal (MIT CAMS) Research Scientist
Panelists: Katrina Hill (Gallagher Re); Kristy Hornland (KPMG); Arjun Ramakrishnan (Mastercard); Bridget Quinn Choi (Gallagher)

12:45PM — Lunch & Member Networking / Poster Session (60 min)

1:45PM — **From Tools to Agents: Cybersecurity in an Era of AI-Driven Attacks** (45 min Industry Keynote) Katie Moussouris, Founder & CEO, Luta Security

2:30PM — Student Research Presentations (30 min Student Presentations) Cynthia Zhang & Lillian Bluestein, MIT CAMS Students Introduced by Dr. Ranjan Pal

3:00PM — Afternoon Break (30 min)

3:30PM — **Securing the Next Generation of Self-Acting Enterprise AI Systems** (60 min Panel)
Moderator: Dr. Michael Siegel (MIT CAMS) Director
Panelists: Ilya Kabanov (MIT CAMS); Katie Moussouris (Luta Security); Sandy Pentland (Stanford/MIT); Paul Quinn (Cisco)

4:30PM — Wrap-Up & Next Steps (15 min Closing Remarks)
Professor Stuart Madnick & Dr. Michael Siegel (MIT CAMS Directors)

5:00PM — Drinks & Networking at Locke Bar (90 min Member Networking

Cybersecurity at MIT Sloan

About CAMS: The Cybersecurity at MIT Sloan (CAMS) team advances research, education, and collaboration across academia, industry, and government. CAMS convenes a global network of executives, scholars, and policymakers to address the evolving challenges of digital security, organizational resilience, and technological change.

Meet the Directors of Cybersecurity at MIT Sloan (CAMS)

Founder & Director Prof. Stuart Madnick

Stuart has been a faculty member at M.I.T. since 1972. He has served as the head of MIT's Information Technologies Group for more than twenty years. During that time the group was consistently rated #1 in the nation among business school information technology programs. He was the Founding Director, and currently co-director, of Cybersecurity at MIT Sloan (formerly the Interdisciplinary Consortium for Improving Critical Infrastructure Cybersecurity.) His involvement in cybersecurity research goes back to 1979 when he co-authored the book Computer Security, one of the first books on this subject.

Co-Founder & Director Dr. Michael Siegel

Michael is a Principal Research Scientist, MIT Sloan School of Management; His research focuses on the management, strategy, technology, and organizational issues related to cybersecurity with specific interest in vulnerability markets, cyber risk management, dark web business models, IoT endpoint security, vulnerability management, cybersecurity workforce development, and educating management in cybersecurity. With over 37 years at MIT, he has published extensively on cybersecurity, AI, and information systems strategy.

Cybersecurity at MIT Sloan

9:15AM — The Escalating Agentic AI Arms Race between the Defenders and Attackers
(45 min Academic Talk)

"As we know AI can help people accomplish their job, even replace them sometimes. What if the job is to be a cyber defender or cyber attacker? What weapons can AI provide?"

Professor Stuart Madnick, MIT CAMS Founder;

Professor of Information Technologies, MIT Sloan School of Management, and Professor of Engineering Systems, MIT School of Engineering

Talk Summary: This presentation introduces many of the key challenges posed by the emergence of Agentic AI in three parts: (1) The Geo-Political Aspect of the Agentic AI Arms Race – Agentic AI is so potential powerful that is fueling an escalating arms race between nations, much like the nuclear arms race, with dangerous consequences; (2) The Defender-Attacker Aspect of the Agentic Arms Race – Both the defenders and attackers of cyber-based systems are using agentic AI as a powerful weapon, sometimes in similar ways and sometimes in ways unique to their position, *for Defenders* it offers the prospect of vulnerability-free code and *for Attackers* it offers a range of new social engineering opportunities; (3) Legal Challenges Created by Agentic AI - In traditional law, agency is when an agent is authorized to act on behalf of a principal legally which binds the principal; but what if the agent does something totally unexpected.

Cybersecurity at MIT Sloan

10:00AM — Offensive AI Is Here. What Do We Do About It?

(45 min Academic Keynote) *"The cybersecurity landscape is undergoing a fundamental shift as adversaries move from simple scripts to autonomous AI agents."*

Dr. Yisroel Mirsky, Zuckerman Faculty Scholar, Head of the Offensive AI Lab, Institute of Software, Systems and Security, Ben-Gurion University

Keynote Summary: The cybersecurity landscape is undergoing a fundamental shift as adversaries move from simple scripts to autonomous AI agents. These systems offer unprecedented speed, scale, persistence, and adaptability, transforming the threat model. At the same time, organizations are adopting agents to scale software development, vulnerability discovery, and security operations—introducing new risks in the very systems meant to improve defense.

In this talk, we examine this transition across four fronts, based on six recent papers from our lab published at venues including Black Hat, USENIX, NDSS, and ICLR. First, the rise of cyber agents: how autonomous agents are changing offensive cyber operations and forcing defenders to rethink control. Second, why code agents are not a silver bullet: how dependence on coding and reasoning agents creates exploitable blind spots in workflows such as malware triage, static analysis, and code review. Third, securing LLM services: how rule-based AI security, model introspection techniques such as PRISM, and agent attribution can help detect, investigate, and hold operators accountable. Finally, we look at what comes next: the spread of these capabilities into criminal ecosystems. Drawing on 141 interviews with insiders in crime syndicates, we examine how AI is beginning to industrialize social engineering and scams at scale.

Together, this talk maps a new era of cyber conflict in which agents operate on both sides—and outlines practical steps we can take now to secure, monitor, and control them.

Cybersecurity at MIT Sloan

11:15AM — When Cyber Risk Becomes Non-linear, Governance Must Become Adaptive

(30 min Academic Talk)

"Through a worse-before-better transition, restoring strategic control"

Dr. Sander Zeijlemaker, CAMS Research Affiliate at Massachusetts Institute of Technology, an agenda contributor to the World Economic Forum, principal research director at Z-CERT, and owner of Disem Institute.

Talk Summary: When Cyber Risk Becomes Non-Linear, Governance Must Become Adaptive

Artificial intelligence is reshaping cybersecurity into a machine-speed competition where attackers and defenders continuously adapt to one another. This creates a "worse-before-better" transition in which offensive AI capabilities initially outpace defensive readiness, increasing systemic cyber risk and exposing the limitations of traditional governance and reactive security models.

In this keynote, we will interactively explore how organizations can restore strategic control and bridge the temporary gap between attacker and defender capabilities. Particular attention will be given to state-sponsored adversaries, whose significant funding, strategic focus, and access to advanced proprietary AI models may further accelerate this imbalance and require other solutions.

Alongside practical short-term actions organizations can implement immediately; the session also looks ahead by introducing CyGENT. This is a consortium-funded research initiative developing solutions for predictive, AI-driven cyber governance. Using data analytics and strategic digital twin technology, CyGENT aims to help leaders anticipate threats, test strategies, and make faster, better-informed decisions in an increasingly non-linear threat environment.

Cybersecurity at MIT Sloan

11:45AM — Poisoned Intelligence: How LLMs Are Reshaping Software Supply Chain Cybersecurity (60 min Panel)

Moderator: Dr. Ranjan Pal (MIT CAMS) Research Scientist

Panelists: Katrina Hill (Gallagher Re); Kristy Hornland (KPMG); Arjun Ramakrishnan (Mastercard); Bridget Quinn Choi (Gallagher)

Panel Moderator

Dr. Ranjan Pal: Research Scientist at MIT Sloan and cybersecurity expert with CAMS. His work focuses broadly on addressing enterprise and supply chain cyber risk and resilience management issues using an interplay of computer science, decision science, and operations research. He is an invited member of the World Economic Forum's (WEF) cyber resilience working group and pioneered quantitative models for cyber insurance market sustainability. Pal has advised corporations and governments on cyber risk management and published more than 100 articles across premier academic and business media outlets. Pal has been an invited speaker in more than 50 venues around the globe and has held visiting scholar and faculty positions at leading academic institutes that include the University of Michigan, the University of Cambridge, Princeton University, King's College London, and Tsinghua University. His work has received multiple best paper award nominations, and he serves as an Associate Editor of the ACM Transactions on Management Information Systems and is a member of the program committees in the Winter Simulation Conference and the ACM SIGSIM-PADS conference.

Panel Members

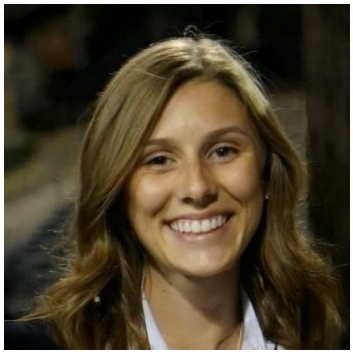

Katrina Hill: is a Senior Cybersecurity Consultant at Gallagher Re, where she leverages her unique cyber expertise to benefit the reinsurance industry. She focuses on translating cyber risks into actionable insights and advising on strategic measures to reduce these risks, ensuring that Gallagher Re clients is better prepared to face the evolving landscape of cyber threats. Prior to Gallagher Re, she was at Accenture where she did six years in Incident Response consulting.

Kristy Hornland: Director, KPMG US. Kristy has deployed Trusted and Secure AI initiatives and governance programs for 20+ Fortune 100 clients while in her role at KPMG US. A recognized industry innovator, she actively transforms client AI adoption by developing comprehensive AI guardrails that integrate leading industry regulations. By deploying robust security strategies and technologies, and sharing insights at forums like Black Hat and SANS, Kristy ensures trusted AI ecosystems while accelerating enterprise innovation. Kristy also leads the AI Security & Trust working group for the Global Resiliency Federation to prepare organizations, public and private, for emerging AI security and trust risk management.

Arjun Ramakrishnan: Arjun Ramakrishnan is a Senior Principal Cybersecurity Architect at Mastercard focusing on Data Protection & AI Security. With over 18 years of diverse experience in information security, his areas of expertise include developing security frameworks and solutions for AI-driven products, with particular emphasis on zero trust, cloud security and identity governance. Arjun is an active participant in academic and industry communities and has contributed to security reports in the field of AI Security. He has spoken at several cybersecurity conferences including the AI Summit, WEF and MIT CAMS, participating in panel discussions alongside leading experts from financial institutions and security organizations Arjun holds a CISM certification, a management degree in Information Systems Management, and an engineering degree in Electronics and Communications.

Bridget Quinn Choi: Senior Vice President, Product Counsel, Cyber and Technology at Gallagher. Bridget has over 20 years' experience in and around the insurance industry. At Gallagher she works with high growth companies and large enterprises to develop insurance solutions for complex cyber, technology, AI, and privacy risks. With a deep understanding of incident response, cyber governance, privacy law, paired with the depth of experience drafting insurance policies and managing claims, she provides coverage solutions and technical services for clients. In her career, she has worked as a cybersecurity consultant, claims counsel, attorney to various companies, and brings a strong understanding of cyber, legal and operational risk.

Cybersecurity at MIT Sloan

1:45PM — From Tools to Agents: Cybersecurity in an Era of AI-Driven Attacks
(45 min Industry Keynote)

Katie Moussouris, Founder and CEO, Luta Security

Keynote Summary: From Tools to Agents: Cybersecurity in an Era of AI-Driven Attacks

AI has crossed a threshold. It is no longer a behind-the-scenes assistant helping analysts spot threats. It is an active participant in cybersecurity itself, capable of reasoning, planning, and acting on its own, for defenders and attackers alike. Defenders are deploying these systems while still writing rulebooks built around human decision-making. Attackers are using the same capabilities to operate at speeds organizations were never designed to withstand. Regulators are drafting policy against a threat landscape that has already moved on.

The hard questions now are about adaptation: what it takes to build defensible organizations against adversaries that no longer play by human constraints, and what accountability looks like when the systems acting on our behalf outpace the institutions meant to govern them. A candid look at where the field actually stands, and what it will take for technologists and society to meet the moment.

Cybersecurity at MIT Sloan

2:30PM — Student Research Presentations (30 min Student Presentations)
Cynthia Zhang & Lillian Bluestein, MIT CAMS Students
Introduced by Dr. Ranjan Pal

Presentation Title: Risk-Driven Strategic Resilience in Digital Infrastructure Networks

Summary: Cyber risk depends not only on how much organizations invest in security, but on which types of capabilities they prioritize. This talk presents a simulation framework that evaluates how investment across the five NIST Cybersecurity Framework functions shapes attacker success in a networked enterprise environment, showing that some strategies materially reduce risk while reactive-only approaches provide little measurable benefit. The results offer a practical basis for prioritizing cybersecurity spending and informing enterprise risk and insurance decisions.

Scan for slides

Cynthia Zhang: is a Master's student in the Electrical Engineering and Computer Science (EECS) department at MIT. She is also a researcher with Cybersecurity at MIT Sloan (CAMS) at the MIT Sloan School of Management. Her primary research interest lies in AI and cyber-security of enterprises/industrial control systems.

Presentation Title: Critical Asset Risk Management for Industrial Control Systems

Summary: This talk presents a network science-oriented approach for prioritizing cyber-risk analysis in industrial control system (ICS) networks. The approach views operational technology (OT) environments as directed, weighted industry equipment networks (graphs) and uses network centrality metrics to rank critical assets, chokepoints, single-points of failure, and high-impact cyber-attack paths. These rankings provide a structured, automated, and cost-effective (w.r.t. time and money) way for SOCs to conduct security analysis and connect technical network structure to practical management decisions around operational risk prioritization.

Scan for slides

Lillian Bluestein: Lillian Bluestein graduated from Massachusetts Institute of Technology with a Bachelor of Science in Artificial Intelligence and Decision Making in May 2026. For the past two years, she has been a student researcher with Cybersecurity at MIT Sloan (CAMS). Her research focuses on strategic cyber risk management for critical infrastructure and business networks, using game theory and graph-based methods to analyze interconnected systems.

Cybersecurity at MIT Sloan

3:30PM — Securing the Next Generation of Self-Acting Enterprise AI Systems
(60 min Panel)

Moderator: Dr. Michael Siegel (MIT CAMS) Director
Panelists: Ilya Kabanov (MIT CAMS); Paul Quinn (Cisco)

Panel Moderator

Dr. Michael Siegel: is a Principal Research Scientist at the MIT Sloan School of Management, He is also the Director of Cybersecurity at MIT Sloan (CAMS). Siegel's research focuses on the management, strategy, technology, and organizational issues related to cybersecurity with specific interest in vulnerability markets, cyber risk management, dark web business models, IoT endpoint security, vulnerability management, cybersecurity workforce development, and educating management in cybersecurity. He also has done research in the intelligent integration of information systems, risk management, insurgency and state stability, data analytics, healthcare systems, and systems modeling. Siegel has published articles on such topics as simulation modeling for cyber resilience, cyber vulnerability markets, A.I. and cybersecurity, data management strategy, architecture for practical metadata integration, heterogeneous database systems, and managing and valuing a corporate IT portfolio using dynamic modeling of software development and maintenance processes. His research at MIT has continued for over 35 years and includes a wide range of publications, patents and teaching accomplishments.

Panel Members

Paul Quinn: is VP of Engineering and Chief Security Architect at Cisco, where he leads platform-wide security strategy across identity, policy, and enterprise infrastructure, including how organizations meet compliance demands without sacrificing the ability to operate securely at scale. A significant focus of his work is AI security, how AI changes the threat landscape and what it takes to deploy AI safely and at scale in enterprise environments. Prior to Cisco, he held senior technical leadership roles at Google Cloud, where he led key initiatives including digital sovereignty.

Ilya Kabanov: is a research affiliate at MIT Sloan Cybersecurity and founder of The Weather Report, a non-profit AI security and safety lab. His research focuses on safe and secure AI. He established and led AI Safety & Security at Google to protect AI products and customers. He also built the global application security and privacy engineering organization at Schneider Electric, the world's leading Industrial IoT company.

He holds an MBA from MIT Sloan School of Management and a Ph.D. from the Moscow Institute of Electronics and Mathematics. He is a senior member of the Institute of Electrical and Electronics Engineers (IEEE).

Sandy Pentland: Alex "Sandy" Pentland is HAI Center Fellow and faculty lead for digital society at Stanford HAI and the Stanford Digital Economy Lab. He is Toshiba Professor at MIT, member of US National Academies, Advisor to Abu Dhabi Investment Authority Lab, and formerly advisory board member at UN Secretary General's office, Google, ATT, Telefonica, and elsewhere. Spin-off companies and open-source systems from his lab manage authentication of most digital transactions in the world, media for roughly 1B people in far east, and health resources for roughly 0.5B people in the indopacific. His current focus is on problems and opportunities in using AI to improve our social institutions.

Katie Moussouris: Founder and CEO, Luta Security. Katie is the founder and CEO of Luta Security, a company that creates and manages sustainable bug bounty and vulnerability disclosure programs. She advises companies and governments on emerging trends in AI and cybersecurity and how the dynamics are shifting between attack, defense,national security, and human expertise adaptation and co-evolution. Ms. Moussouris is the co-author and co-editor of ISO standards 29147 vulnerability disclosure and ISO 30111 vulnerability handling processes. During her tenure with Microsoft, Ms. Moussouris started Microsoft Vulnerability Research as well as Microsoft's first bug bounty program. She served as a technical advisor to the US Wassenaar delegation to help renegotiate clarifications to export control of intrusion software and technology. Ms. Moussouris led the launch of the U.S. government's first bug bounty program, "Hack the Pentagon."

Cybersecurity at MIT Sloan

Current Research Projects

Risk-Driven Strategic Resilience in Digital Infrastructure Networks

Abstract: Cyber risk depends not only on how much organizations invest in security, but on which types of capabilities they prioritize. This research presents a simulation framework that evaluates how investment across the five NIST Cybersecurity Framework functions shapes attacker success in a networked enterprise environment, showing that some strategies materially reduce risk while reactive-only approaches provide little measurable benefit. The results offer a practical basis for prioritizing cybersecurity spending and informing enterprise risk and insurance decisions.

Project Leads: Ranjan Pal, Michael Siegel, Gabriel Bassett (Liberty Mutual Insurance)

Managing Cyber Risk in AI and Software Supply Chain Networks

Abstract: Most enterprise AI and software applications (AISAs) are built atop components (i.e., code, data, and models) that are either open sourced or from third parties and form a dependency network. Consequently, they are exposed to cyber vulnerabilities of varying criticalities. While it is imperative for enterprise AISA products to be security vetted before deployment, it is not practically or computationally feasible for cyber risk management to address every weak link in the component dependency network that underlies such products. We develop a graph-theoretic methodology validated using real world datasets and Monte Carlo simulations to identify open-source and third-party AISA components that have the maximum impact on the security of enterprise supply chains reliant on such enterprise AISA products. We subsequently model and rank using graph theory cybersecurity KPI metrics for effectively governing supply chain cybersecurity.

Project Leads: Ranjan Pal, Michael Siegel, KPMG Team

Scaling Cyber Insurance Markets Using Financial Catastrophe (CAT) Bonds

Abstract: Cyber insurance markets are inefficient and struggle to scale due to amplified correlated risks from tightly interconnected IT and operational technology (OT) ecosystems and weak cyber vulnerability disclosure rules. However, these markets improve the cybersecurity of companies (enterprises) in theory. This theory is increasingly being validated in practice in multiple enterprises that buy cyber insurance. However, the supply-demand gap in the CI market is huge and stretches over a trillion dollars. The primary reason being that cyber reinsurance companies do not have sufficient capital to serve their CI company clients through profitable cyber-risk portfolio diversification. This capital problem is subsequently rooted in the fundamental challenge of enterprise cyber posture information asymmetry that has plagued the CI industry since its inception. A radical capital-boosting mechanism proposed by researchers and deployed within the industry in the last two years is an insurance-linked security (ILS) product such as a catastrophic (CAT) bond. In this project, we present arguments on *why*, *how*, and *when* CAT bonds

help sustain capital-boosted reinsurance markets complemented by innovative AI/ML-driven inside-out cyber posture scanning solutions. We append the theory behind these arguments with an AI-based Monte Carlo simulation theory that learns from real-world enterprise loss and posture data statistics, solves for market equilibria, and tightly bounds market KPIs. Our key result shows that CAT bonds can effectively resolve the scaling problem in the cyber insurance industry only for manageable costs of mitigating information asymmetry (IA) by CAT bond investor brokers on enterprise cyber posture.

Project Leads: Ranjan Pal, Michael Siegel, Stuart Madnick

Rethinking Insurance for the AI and Agentic AI era

Abstract: The increase of artificial intelligence (AI) and agentic AI systems across enterprise operations has created a fundamental transformation in organizational risk profiles. While AI and its agentic counterpart deliver measurable operational efficiencies and competitive advantages, it simultaneously introduces novel vulnerabilities, expands attack surfaces, and challenges traditional security frameworks. This dual nature positions AI as both strategic asset and potential liability, creating urgent questions about appropriate risk transfer mechanisms. Current insurance products, such as cyber, errors and omissions, product liability, and motor insurance, provide fragmentary coverage for AI-specific risks. The insurance industry faces a fundamental debate: should AI and agentic exposures be addressed through adaptation of existing policies or through entirely new insurance products? This question is complicated by ambiguous policy language, limited actuarial data, and the technical complexity of AI and agentic systems that straddle multiple traditional coverage boundaries. Liability allocation between AI developers and AI deployers further complicates coverage design. Recent litigation demonstrates that deploying organizations bear disproportionate risk burden, facing legal and financial consequences for AI system failures even when vendors created the underlying technology. This asymmetry creates significant gaps in coverage where operational losses occur without clear insurance pathways. This project examines systematic coverage gaps in current insurance policies, analyses emerging AI and agentic specific risks through documented incidents, and evaluates both standalone products and enhanced endorsements as potential solutions.

Project Leads: Ranjan Pal, Katrina Hill (Gallagher Re)

The Theory and Practice of Managing Systemic Risk in Business Ecosystems

Abstract: Business enterprises have grappled in the last one and half decade with unavoidable risks of (major) cyber incidents on critical infrastructure (e.g., power grid, cloud systems). The market to manage such risks using cyber insurance (CI) has been growing steadily (but not fast enough) as it is still skeptical of the extent of economic and societal impact of systemic risk across networked supply chains in interdependent IT-driven enterprises. To demystify this skepticism, we first study in this project the role of (a) the statistical nature of multiple enterprise cyber-risks contributing to aggregate supply chain risk and (b) the graph structure of the underlying enterprise supply chain network, in the statistical spread of aggregate cyber-risk. More specifically, we provide statistical tail bounds on the aggregate cyber-risk that a cyber-risk management firm such as a cyber insurer is exposed to in a supply chain. Subsequently, we study the problem of aggregate cyber-risk management by cyber re-insurance firms via portfolio

design to optimally diversify aggregate/systemic cyber-risk sourced from multiple CIs insuring enterprises on a supply chain. As a first, we propose a mathematical framework for cyber re-insurers to test the operational sustainability of systemic cyber-risk diversification portfolios with respect to the standard Value-at-Risk (VaR) metric for general aggregate cyber risk distributions. In addition, we propose a statistical copula methodology to make systemic cyber-risk portfolio diversification sustainable for cyber re-insurers in scenarios where the sustainability test fails. We validate our theory via Monte Carlo simulations.

Project Leads: Ranjan Pal, Michael Siegel

Navigating AI and Global Policy: Artificial Intelligence, Cybercrime, and the Emerging Regulatory Landscape

Abstract: Artificial intelligence is reshaping cybercrime at unprecedented speed. AI-driven attacks — from deepfakes and synthetic identity fraud to autonomous phishing — are becoming more scalable, adaptive, and accessible, empowering even non-expert actors to cause significant harm. Yet the same technologies also enhance detection, prevention, and forensic analysis, creating a profound regulatory paradox that existing legal frameworks remain poorly equipped to manage. This project examines how global and regional policy frameworks govern the dual-use nature of AI in cybercrime, and where critical gaps persist. Grounded in an empirical analysis of AI-enabled cybercrime incidents, it shows that identity-based attacks account for more than half of documented cases, that harms are multidimensional, and that privacy risks are systemic rather than exceptional. A central contribution is the articulation of an "attribution crisis": as AI systems act with growing autonomy across jurisdictions, legal frameworks struggle to assign responsibility and criminal liability. Through comparative analysis of national, regional, and international instruments — including the EU AI Act, selected U.S. state legislation, and the 2024 UN Convention against Cybercrime — the project maps where current provisions apply, where they break down, and what organizational and technical preparedness measures are necessary for compliance readiness. It also expands the comparative scope to consider additional jurisdictions, such as Japan, as part of a broader effort to assess how different legal systems are responding to AI-enabled cybercrime. In doing so, it offers a timely framework for policymakers and organizations confronting one of the most urgent governance challenges in the evolving relationship between AI and cybersecurity.

Project Leads: Angelica Marotta, Stuart Madnick.

Automated Critical Asset Cyber Risk Management in Industrial Control Systems

Abstract: Critical assets in industrial control systems (e.g., PLCs, historians, HMIs, safety controllers) are essential components whose compromise can disrupt operations, personnel safety, or data integrity across interconnected OT environments. The scale of such sensor and actuator networked environments is often quite large within industrial control systems (ICSs). The popularly known primary critical asset risk management (CARM) challenge faced by most IT/OT cybersecurity managers globally is the lack of a principled and automated CARM framework to improve ICS performance metrics. In this research, we develop a principled graph-theoretic methodology validated using real world trace driven Monte Carlo simulations to identify the most important ICS critical assets that have the maximum adverse impact on

ICS performance if compromised. We model and rank using graph theory cybersecurity KPI metrics of *blast radius* and *single points of failure* (SPOFs) for ICS network assets for effectively governing ICS CARM and cybersecurity.

Project Leads: Ranjan Pal, Michael Siegel, IPA Team

Strategic Cyber Resilience of Flow Applications in Industrial Control Systems

Abstract: Across industrial control system (ICS) networks such as power and transportation networks, ``flow" applications are the optimization-and-control loop applications that balances supply against demand under system constraints. These constraints include economically balancing load among network components, ensuring system variables are within operational limits, and managing network congestion. It is of growing interest to geo-political cyber adversaries today that they disrupt flow applications in digitally driven critical infrastructure system networks such as in ICSs along strategic links of the network. Doing so would result in basic civilian service disruption (e.g., loss of power, lack of potable drinking water) and cause the society and economy significant harm. Add to this the obvious fact that it is impossible to always defend ICS networks from highly resourced and strategic adversaries. Hence, effective resilience to ICS network flow cyber-attacks rather than their sole prevention has become a priority for critical infrastructure management. Given an ICS network supporting a flow application at a standard QoS, we are interested in optimizing cyber resilience of the latter post an inevitable cyber-attack that reduces this QoS. Here, resilience is the quantitative ability to anticipate and respond post a cyber-attack and ensure that standard flow QoS is re-instated within a time constraint. We formally model this time-dependent cyber resilience optimization problem as a mixed-integer linear program (MILP) - the solution to which enables ICS management to schedule cyber-defense operations on adversary-affected network edges in a deploy-friendly manner that optimizes resilience. The set of affected edges is determined by the competitive outcome between a strategic adversary and a proactive strategic defender akin to the seminal *Tullock* contest in game theory and one that is applicable to stealthy APT-type cyber threats for flow disruption. We model adversary and defender strategies in this contest that are both AI-based and non AI-based with respect to targeting flow edges central to the ICS network. We prove the computational complexity of the optimization problem and study the comparative effect of multiple adversary-defender strategy combinations on the optimal value of flow application cyber resilience. We conduct extensive Monte Carlo simulations of our proposed network resilience optimization approach on the standard and real-world IEEE 30-bus and IEEE 118-bus architectures in power system ICSs.

Project Leads: Ranjan Pal, Michael Siegel

An Economics of Sharing Cyber Threat Intelligence in Business Networks

Abstract: Digital supply chain networks (DSCNs) are expanding as businesses embrace IT, AI, and IoT/CPS, collectively widening the DSCN attack surface. Experts urge cyber threat intelligence (CTI) sharing by DSCN enterprises, yet 'toothless' regulations and competitive fears over exposing confidential data hamper collaboration towards realizing trustworthy DSCNs. We introduce a game-theoretic framework for strategic CTI sharing across DSCNs. Using a scalar metric acting as the CTI quality

measure, our economics model is the first to concurrently captures network topology, spillover effects, quality-privacy tradeoffs, and decentralized competition among CTI sharing DSCN firms. Theory and Monte Carlo simulations show that enterprise *Katz-Bonacich* (KB) centrality within a DSCN together with the strategic complementary nature of its shared CTI dictate (i) its optimal CTI sharing quality at a Nash equilibrium (NE) and (ii) the DSCN's *price of anarchy* (PoA) at the NE that denotes the additional sharing quality improvement collectively possible for strategic DSCN enterprises.

Project Leads: Ranjan Pal, Michael Siegel

Cyber Defense Frameworks to Manage LLM and Agentic AI Supply Chain Risks

Abstract: AI and agentic supply chains are now part of core digital infrastructure: models, data, tools, connectors, and agent runtimes are being deployed into production workflows that touch sensitive data and real-world actions. Unlike traditional software, failures can be silent, scalable, and self-amplifying; poisoned training data becomes "policy," a compromised artifact or registry can spread fleet-wide, and agent toolchains can turn untrusted context into privileged execution. Joint research with KPMG and MIT CAMS report that the dominant risks shift from single-component vulnerabilities to end-to-end lifecycle and runtime supply-chain exposure, including provenance gaps, update channels, tool/memory poisoning, identity delegation, and human oversight. Managing such risks is essential to reduce systemic business risk, support trustworthy adoption, and give leadership a practical roadmap: what to sign, gate, monitor, and govern before the organization ships autonomy at scale. A joint report on LLM supply chain risk management frames risk across the full lifecycle - data/toolchains, development, release/ops, distribution, and people - where the biggest failures come from poisoned inputs, compromised build/update paths, data extraction, endpoint/API abuse, and insider/vendor misuse. It emphasizes leadership priorities like proving provenance (DBOM/MBOM, dataset lineage, signed artifacts, pinned dependencies), gating promotion (secret evals, red-teaming, tiered releases/approvals), and monitoring/containment (telemetry, prompt-injection defenses, exfil detection, budget/rate circuit breakers). A joint report on agentic AI supply chain risk management extends this to runtime autonomy, where the "dependency graph" includes tools (MCP), memory/RAG, identity/delegation, multi-agent orchestration, and human oversight. It highlights attack classes like tool poisoning/rug-pull, indirect prompt injection and memory poisoning, over-privileged non-human identities, and multi-agent cascade/fleet compromise, and recommends: sign provenance end-to-end, govern agents as NHIs (per-task identities, OAuth on-behalf-of, prompt lineage), and place a supervisor on every boundary (tool calls and agent-to-agent hops) with budgets and human approval gates.

Project Leads: Ranjan Pal and KPMG team

An Agentic-Synthetic Framework for Multimodal Phishing Risk Management

Abstract: Modern cybersecurity programs increasingly depend on AI for multimodal (e.g., email, audio) phishing detection across IT environments, yet the core input to these systems - high-quality labeled phishing data is structurally scarce. Real multimodal phishing data is proprietary, expensive to label, and inherently rare (especially in legacy and critical-infrastructure systems), while "normal" behavior is domain-specific and shifts across enterprises, causing pre-trained models to fail when transferred. This motivates

a central research problem: *How can we systematically leverage small pools of real, shareable seed data to generate agentic-synthetic multimodal phishing datasets that are sufficiently realistic and diverse to train accurate phishing detectors, while providing quantitative guarantees on both synthetic-data fidelity and downstream detection performance?* Concretely, the research problem requires designing controlled agentic AI experiments that amplify small seed datasets, validate synthetic realism with statistical metrics, and run Monte Carlo sweeps across seeds, scaling, and models to quantify confidence bounds; performing controlled test on human subjects for synthetic phishing object effectiveness; then translate results into board KPIs like number of alerts, precision/recall, SOC cost, avoided loss, and net benefit.

Project Leads: Ranjan Pal, Michael Siegel

From Institutional Contradictions to Disturbances: Understanding Cybersecurity Resilience in Electricity Infrastructure

Abstract: Cybersecurity threats targeting electricity infrastructure (EI) have intensified in recent years, and despite substantial investments and increasingly stringent regulations, vulnerabilities remain persistent and systemic. Existing research has largely treated these challenges as technical, organizational, or regulatory deficiencies, thereby offering limited explanation as to why they endure across contexts and over time. To address this gap, we conceptualize these challenges as manifestations of deeper institutional contradictions rather than isolated failures. Using a dialectical analysis of 30 interviews with diverse EI stakeholders in the U.S. and Finland, we identify six recurrent dialectical poles that underlie cybersecurity disturbances in EI operations. These poles emerge from enduring tensions between efficiency, legitimacy, flexibility, stability, and resilience rooted in the distinctive socio-technical and institutional characteristics of EI. Our findings offer a novel explanation for the persistence of cybersecurity vulnerabilities and provide a foundation for interventions that address root institutional causes rather than symptomatic technical failures.

Project Leads: Bui, Vartiainen, Proudfoot, Dang, Madnick, Coden

Recommendations for Advancing Cybersecurity Governance in the Boardroom to Address Persistent and Emerging Challenges

Abstract: Cyber risk has emerged as a particularly urgent concern due to its potential for existential disruption and mounting regulatory pressures. Despite corporate boards' increased attention to cyber risk, old challenges persist as new threats emerge. Our 27 hours of interviews with 42 board directors, board advisors and board-facing executives identified two core types of governance challenges: (1) internal structural issues and (2) external cyber threats. For each challenge, we offer expert-backed guidance on how to anticipate, prepare and respond. We conclude with challenges we believe must still be addressed.

Project Leads: Proudfoot, Jeffrey G.; Cram, W. Alec; Madnick, Stuart; and Coden, Michael

An AI–Cybersecurity Taxonomy for Regulating Hybrid Risks

Abstract: This project responds to an urgent regulatory challenge: the growing convergence of artificial intelligence and cybersecurity risks is outpacing governance models that still treat these domains as separate policy fields. As AI systems become increasingly embedded in critical infrastructures, digital services, and security operations, regulators and institutions lack a coherent framework for identifying and managing the hybrid risks that arise at their intersection. The project develops the AI–Cybersecurity Taxonomy, a structured governance framework that maps five core regulatory clusters through which these converging risks can be analyzed and addressed. Drawing on comparative regulatory analysis and sector-specific case studies, it examines how existing frameworks are adapting—or failing to adapt—to issues such as agentic amplification, embedded verifiability, and provisional governance. In doing so, it provides a more integrated conceptual vocabulary for policymakers, compliance professionals, and organizational leaders working across AI and cybersecurity domains. Key outcomes include a taxonomy for classifying hybrid risks, a set of analytical concepts for understanding emergent governance challenges, and practical recommendations for more flexible, layered regulatory design. By advancing an integrated approach to AI and cybersecurity oversight, the project supports more resilient, adaptive, and policy-relevant governance strategies.

Project Leads: Angelica Marotta, Stuart Madnick.

The Insider Threat You Can't Fire: Cybersecurity Whistleblowers

Abstract: Cybersecurity whistleblowers operate where institutional silence becomes systemic risk. Often the first to detect hidden failures, governance breakdowns, and unreported vulnerabilities, they can determine whether organizations confront cyber threats early or allow them to escalate into crisis. This project examines how insider disclosures reshape cybersecurity from a technical issue into a decisive question of law, governance, and organizational resilience. Through analysis of legal frameworks, enforcement actions, and reporting practices across public and private institutions, it identifies when whistleblowing functions as a credible early-warning mechanism and when it is ignored, suppressed, or rendered ineffective. The project highlights the conditions that transform internal reporting into actionable intelligence, strengthen protection against retaliation, and improve organizational response to emerging threats. In doing so, it positions whistleblowers not as marginal actors, but as pivotal agents of transparency, prevention, and institutional accountability in contemporary cybersecurity governance.

Project Leads: Angelica Marotta, Stuart Madnick.

Cybersecurity Metrics for Boards of Directors

Abstract: Cybersecurity oversight is increasingly central to boards of directors due to regulatory pressures and the growing recognition of cyber risk as a major organizational threat. Yet boards lack practical metrics to support effective governance. Existing measures often focus on board composition or rely on overly technical indicators that offer limited decision-making value. This article introduces a novel set of cybersecurity metrics designed for board-level use. Grounded in financial accounting literacy, the metrics translate cyber risk into familiar financial terms. Developed and evaluated through an industry-focused

approach, they support better governance decisions, inform cybersecurity investments, improve regulatory reporting after breaches, and help boards assess organizational risk exposure and strengthen overall cybersecurity oversight.

Project Leads: Schaffner and Proudfoot

Drowning in Compliance, Starving for Security: The Cybersecurity Regulation Paradox

Abstract: The increasingly hostile and dangerous digital landscape is resulting in frequent and severe cyberattacks in virtually every industry. This creates a mounting pressure on policymakers who intervene with the intention to protect data, ensure consumer privacy, reduce cyber risk, and at a more macro level, shore up national security. Policymakers' response to this pressure often takes the form of cybersecurity regulations that compel organizations to comply with prescriptive requirements. Non-compliance is penalized with financial fines and/or other punitive measures. However, researchers have yet to conclusively understand the nuances and effects of cybersecurity regulation development, dissemination, and organizational response. In this research, we explore a subset of this problem space: the potential deficiencies of cybersecurity regulation development and dissemination. To do so, we chose a qualitative approach and conducted 22 semi-structured interviews with a heterogeneous sample of cybersecurity and regulatory experts. Sensitised by paradox theory and using an inductive coding analysis, we identified three core tensions that reflect the deficiencies emerging from our data associated with the development and dissemination of cybersecurity regulations, which, counterintuitively, may harm the cybersecurity posture of compliant organizations. We discuss these deficiencies and provide illustrative quotes to solidify their salience. We conclude with the implications and limitations of our work and propose avenues for future research.

Project Leads: Proudfoot, Demetis, and Madnick

Understanding Cybersecurity Resilience in Electricity Infrastructure: A Configurational Analysis of Challenges and Solutions

Abstract: As electricity infrastructure (EI) undergoes rapid digital transformation, organizations face increasing cybersecurity risks stemming from heightened system interconnectivity, expanding attack surfaces, geopolitical instability, domestic threats, and climate-related disruptions. Despite the critical importance of cybersecurity resilience for national security and societal stability, existing Information Systems (IS) research has paid limited attention to how combinations of cybersecurity challenges and solutions collectively shape resilience in the EI sector. Addressing this gap, this study adopts a configurational perspective to examine how distinct combinations of organizational, technological, and environmental factors influence cybersecurity resilience in EI organizations. Using a mixed-methods research design, Phase 1 consists of 20 qualitative interviews with stakeholders across the EI supply chain to identify key cybersecurity challenges and mitigation strategies. Building on these findings, Phase 2 employs Qualitative Comparative Analysis (QCA) to identify configurations of cybersecurity challenges and solutions associated with varying levels of cybersecurity resilience. By advancing a holistic and

configurational understanding of cybersecurity resilience, this research contributes to the growing literature on energy informatics, digital sustainability, and critical infrastructure cybersecurity.

Project Leads: Bui, Vartiainen, Proudfoot, Dang, Madnick, Coden

Intelligence, Information Operations, and Cyber Physical Security in the Middle East, with a Focus on Saudi Arabia

Abstract: This project examines how artificial intelligence is increasingly reshaping intelligence, information operations, and cyber physical security in national security competition around the world. As AI enabled technologies are integrated into IoT platforms, industrial control systems, and smart city infrastructure, they are changing how information is collected, analyzed, and acted upon across civilian and strategic domains. These capabilities expand analytic and operational capacity, but they also create new risks related to governance, resilience, deception, disruption, and cascading failure. Using Saudi Arabia as a case study, the project analyzes how oversight of AI enabled systems is distributed across cybersecurity authorities, data governance institutions, sector regulators, infrastructure operators, and technology vendors. Institutions such as the National Cybersecurity Authority and the Saudi Data and Artificial Intelligence Authority provide strategic direction through cybersecurity controls and AI governance frameworks. Yet questions remain regarding cross sector coordination, accountability for AI enabled decision making, and dependence on global technology supply chains. By tracing Saudi Arabia's evolving regulatory landscape, the project argues that governance ambiguity can become a source of systemic risk as AI becomes more deeply embedded in intelligence functions and critical infrastructure.

Project Leads: May Almousa, Stuart Madnick

Synthetic Cybersecurity Data for AI Based Threat Detection

Abstract: This project addresses one of the central challenges in AI based cybersecurity: the scarcity of high quality, labeled, shareable cyber data. Real cyber incident data is often proprietary, sensitive, incomplete, or difficult to label, limiting the development and evaluation of detection models. The project investigates how synthetic data generation can be used to expand small real-world datasets while preserving statistical realism and improving downstream detection performance. Potential use cases include phishing, malware, ransomware, OT and ICS logs, and enterprise security alerts.

Project Leads: May Almousa, Ranjan Pal, Michael Siegel

Cygent Project

MIT CAMS, University of Twente and Disem Institute have been awarded a 0.7 mln. EUR grand (~ 850 K US$). The Cyber Governance and AI-Driven Threat Management project (CyGENT) focuses on the growing role of AI within the malware ecosystem and its impact on cyber threats and their management. By leveraging large-scale threat data sharing and transparency tools such as VirusTotal and Malware

Bazaar, the project aims to improve malware classification by distinguishing AI-driven threats from non-AI-driven ones.

Project Leads: Sander Zeijlemaker, Michael Siegel, May Almousa

Reportable or Not? Determining Cyber Incident Reportability Under Uncertainty

Abstract: This project addresses a critical governance gap in cyber incident response: how organizations can determine whether a cybersecurity incident is reportable when evidence remains incomplete, ambiguous, and time-sensitive. As regulatory obligations expand across jurisdictions, institutions face increasing pressure to make defensible disclosure decisions under conditions of uncertainty, often without a structured method for aligning technical facts with legal reporting thresholds. The project develops a practical, comparative decision framework to support timely and accountable reporting determinations across four major regulatory regimes—GDPR, NIS2, SEC cybersecurity disclosure, and HIPAA. Using comparative legal and governance analysis, it identifies and translates differences in harm models, trigger conditions, and reporting timelines into a common multi-lens structure for institutional decision-making. The framework draws on public cases, such as the Conduent's 2025 cyber incident, to demonstrate how reporting obligations can emerge asynchronously as facts evolve. A central output of the project is a tri-state model that distinguishes between incidents that are below threshold, not yet evidenced, or already reportable, along with a defensible approach for documenting uncertainty and escalation. By improving clarity at the intersection of law, governance, and operational response, the project strengthens cyber incident accountability, organizational coordination, and regulatory compliance.

Project Leads: Angelica Marotta, Stuart Madnick.

A Systems Dynamics Approach to Understand Threat Intelligence Sharing.

Abstract: Cyber threat intelligence (CTI) sharing is widely regarded as essential for collective cyber defense; however, its effectiveness depends on fragile institutional incentives and timing, particularly concerning zero-day vulnerabilities. This paper conceptualizes CTI sharing as a dynamic interaction between attackers and defenders, shaped by information asymmetry, disclosure, and mitigation delays. Using a system dynamics approach, we develop and validate a state-based model of vulnerability lifecycles by reproducing empirical zero-day survival curves and historical *Common Vulnerabilities and Exposures (CVE)* disclosure patterns 0f two different datasets. Simulation results indicate that weakening centralized disclosure mechanisms significantly increases societal cyber risk, whereas improvements in secure-by-design practices, accelerated patching, and stronger disclosure incentives reduce vulnerability accumulation. The findings underscore CTI sharing as a critical public good and emphasize the dual role of artificial intelligence in amplifying both software risk and defensive capacity.

Project Leads: Sander Zeijlemaker, Ranjan Pal, Michael Siegel

Multi-Agent Adversarial Architectures and Cyber Risk Dynamics in Critical Infrastructure Environments

Abstract: Artificial intelligence is enabling autonomous multi-agent systems that can coordinate cyberattacks at machine speed, creating new asymmetries between attackers and defenders. This research investigates how swarm-based adversarial architectures influence resilience, capability degradation, and cyber risk in critical infrastructure environments. The study follows a three-stage approach. First, adversarial attack and defense architectures are developed and explored to understand coordination, scalability, and attacker–defender dynamics. Second, insights from these architectures are translated into system dynamics models to analyze capability erosion, escalation behavior, resilience, and governance interventions. Third, optional, experimental validation deepens the research through cyber and cyber-physical testing environments, strengthening strategic and operational findings.

Project Leads: Sander Zeijlemaker, Michael Siegel and Rishi Koduri

Recent and Imminent Publications

Marotta, A., & Madnick, S. (2026). An AI–Cybersecurity Taxonomy for Regulating Hybrid Risks [Manuscript submitted to the International Association for Computer Information Systems 2026 conference].

Marotta, A., & Madnick, S. (2026). Reportable or not? Determining cyber incident reportability under uncertainty [Manuscript submitted to the International Conference on Information Systems 2026 conference].

Marotta, A., & Madnick, S. (in press). A systematic understanding of regulatory definitions of "cyber incidents". Journal of Information System Security [accepted for publication].

Marotta, A., & Madnick, S. (2026). A feature-driven analysis of global cybersecurity regulations and their impact on safeguarding data value. Journal of Data and Information Quality, 18(2), Article 591.

Proudfoot, J., Cram, W. A., Madnick, S., Coden, M. (June 2026). Cybersecurity Governance in the Boardroom: Persistent and Emergent Challenges with Recommendations for Action. Management Information Systems Quarterly Executive (MISQE) (25) 2.

Bui, Q., Proudfoot, J. G., Adelakun, O. (2026). IT Governance in the Boardroom: The Emergence of Boards as Digital Fiduciaries in the Information Age. Proceedings of European Conference on Information Systems (ECIS). Milan, Italy.

Proudfoot, J. G., Bui, Q., Madnick, S., Coden, M. (2026). "Cybersecurity Challenges in Electricity Infrastructure: A Socio-Technical Perspective" Proceedings of 25th Annual Security Conference. Las Vegas, Nevada.

Tony Delvecchio, Sander Zeijlemaker, Giancarlo De Bernardis, Michael Siegel, Human-centered interface design for a dynamic cyber-risk group-based training game, Computer Standards & Interfaces, Volume 95, 2026, 104030, ISSN 0920-5489, https://doi.org/10.1016/j.csi.2025.104030

Zeijlemaker, Siegel, *How to prioritize cyber resilience in the healthcare sector*. Wold Economic Forum, https://www.weforum.org/stories/2026/02/how-to-prioritize-cyber-resilience-in-the-healthcare-sector

Delvecchio, T., Zeijlemaker, S., De Bernardis, G., Siegel, M. (2026). Employing Board Cyber-Risk Management Collaborative Game Under Condition of Uncertainty. In: Lenzini, G., Mori, P., Furnell, S. (eds) Information Systems Security and Privacy. ICISSP ICISSP 2023 2024. Communications in Computer and Information Science, vol 2459. Springer, Cham. https://doi.org/10.1007/978-3-031-89518-0_3

Recent Press

S. Madnick and J. Proudfoot: Boards Are Falling Short on Cybersecurity. Harvard Business Review, April 2, 2026.

R. Pal: Smart Systems, Blind Spots: Rethinking Insurance for the AI Era. Gallagher Re Report, March 25, 2026.

R. Pal et al.: The Rising APT Risk: Reshaping Cyber Insurance for Critical Infrastructure. Forbes India, March 17, 2026.

S. Madnick and A. Marotta: What the UN Treaty on Cybercrime May Mean for You. MIT Sloan Management Review, February 9, 2026.

R. Pal, S. Zeijlemaker, and B. Nag: Synthetic Data: The New Backbone of Next-Gen Cybersecurity. Forbes India, February 5, 2026.

S. Zeijlemaker and M. Siegel: How to Prioritize Cyber Resilience in the Healthcare Sector. World Economic Forum, February 4, 2026.

R. Pal et al.: Cyber Risk in the Boardroom: Why Judgment Matters More Than Numbers. Forbes. January 19, 2026.

R. Pal, A. Tuteja, and B. Nag: Businesses Need Future-Ready LLM Supply Chains. Forbes India, December 30, 2025.

Zeijlemaker, S., Siegel, M., & Gavgani, B. (2025) *IA et gestion des risques cyber: les quatre tendances émergentes à surveiller*. Alliancy. English Translation Here.

R. Pal, D. Yao, and B. Nag: How Can Companies Secure Their Future with Rise in Agentic AI Adoption. Forbes (I), October 30, 2025.

R. Pal, B. Nag, and S. Mukherjee: Cyber Risk on India's Electric Roads. Forbes (I), October 16, 2025.

R. Pal and B. Nag: Quantum Risk, National Fallout: Rethinking Cybersecurity Before its Too Late. Forbes (I), September 30, 2025.

R. Pal "Future-proofing your company from quantum cyber risks," Forbes (I), December 12,2025.

R. Pal, D. Granillo, and B. Nag: How AI Hallucinations Endanger Software Supply Chain. Forbes (I), August 28, 2025.

Marotta, Angelica, & Madnick, Stuart. (2025). The UN Cybercrime Treaty and AI: Navigating the intersection of technology and global policy. Issues in Information Systems, 26(4), 1-16.

Marotta, Angelica, & Madnick, Stuart. (2025). A comparative analysis of the UN Cybercrime Treaty and GDPR: Balancing global security and data sovereignty. Proceedings of AMCIS 2025 – SIG_SEC.

Marotta, Angelica, & Madnick, Stuart. (2025). Analyzing and categorizing emerging cybersecurity ACM Computing Surveys.

Y. Lemiesa, R. Pal, and M. Siegel: *AI on Small and Noisy Data is Ineffective for ICS Cyber Risk Management.* To Appear in Proceedings of Winter Simulation Conference [an INFORMS event], 2025 (IEEE Press), Seattle, Washington, USA. ['Best Paper Award Finalist' in the *Theory* track]

R. Pal, K. Duan, and R. Sequeira: *A Theory to Estimate, Bound, and Manage Systemic Cyber Risk.* In Proceedings of ACM SIGSIM PADS, 2025, New Mexico, USA.