
The Science of Cyber Risk: A Research Agenda

By Gregory Falco¹, Martin Eling, Danielle Jablanski, Virginia Miller, Lawrence A. Gordon, Shaun Shuxun Wang, Joan Schmit, Russell Thomas, Mauro Elvedi, Thomas Maillart, Emy Donavan, Simon Dejung, Matthias Weber, Eric Durand, Franklin Nutter, Uzi Scheffer, Gil Arazi, Gilbert Ohana, Herbert Lin

Cyber risk encompasses a broad spectrum of risks to digital systems, such as data breaches or full-fledged cyber attacks on the electric grid. Efforts to systematically advance the science of cyber risk must draw upon not only computer science, but also fields such as behavioral science, economics, law, management science, and political science. Yet many scholars believe that they have sufficient understanding of other fields to comprehensively address the inherently cross-disciplinary nature of cyber risk. For example, a statistician might apply Bayesian modeling to predict future cyber events, even though it is not entirely clear what bearing historical cyber events have on future ones. Computer scientists might write on data protection laws, yet with little knowledge of legal jurisdiction issues. Such questions of disciplinary ownership, the inability to coordinate across disciplines, and the undefined scope of the problem domain have thus plagued inherently cross-disciplinary cyber risk research. Drawing upon global expertise and challenges from industry, academia, non-profit organizations, and governments, we adapt the classical risk management process to identify core research questions for cyber risk, gaps in knowledge that need to be addressed for advances in security, and opportunities for cross-disciplinary collaboration for each area. While we mention specific disciplines, reflective of our backgrounds, these are not the only ones that should be conducting cyber risk research.

WHAT CONSTITUTES CYBER RISK?

We consider cyber risks to include “operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems” (1). But the scope of cyber risk has nonetheless been difficult to characterize because there is disagreement about what a cyber event and other pertinent terminology actually entail. Some organizations consider a cyber event to be any unknown connection attempt to their network, but others only consider successful unauthorized access to their network. Still others only consider instances when “loss” is experienced. Such definitional ambiguity is problematic, making it difficult to benchmark security or risk across an industry or even within the same organization. For the purposes of establishing a body of knowledge, clear and consistent terminology is essential (2).

Recently there has been welcome progress. For example, the International Organization for Standardization established cyber guidance for insurance and machinery. But as seen with the U.S. government’s National Institute of Standards and Technology’s cybersecurity framework, implementation has faltered among those seeking to employ the “gold standard” of cyber risk management. One important reason is the high costs associated with appropriate security controls and the lack of qualified people to implement the standard. In our view, in addition to increased convergence on cyber guidance and standards, there is need for a standard of care that defines how an organization can proceed with enacting a standard. By addressing cyber risk terminology, standards, and implementation principles in a cross-disciplinary fashion, such guidance can be interpretable and usable by a wide variety of companies, non-profit, and governmental organizations that have different agendas. Diversity of thought will likely contribute to richer cyber risk insights.

There is also value to classifying the variety of cyber risks so that they can be addressed appropriately. For example, cyber risks can be classified by motivation (e.g. financial, reputational), type of attack (e.g. malware, insider attack, spam, distributed denial of service) and source (e.g. terrorists, criminals, government). Clarification on types of cyber risk can help to develop a common language for the science. Progress has been made by the non-profit MITRE Corporation in classifying risk types and establishing terminology. But adaptations of classical risk management tools like early detection systems or risk maps are still in their infant stages, also because of measurement and quantification difficulties.

An example of disciplines that could benefit from collaborating on this topic is computer science and law. Increasingly legislation is being proposed about technical attributes of cyber risk. Computer scientists should contribute to proposed technical regulatory requirements and definitional boundaries. Without collaboration here, legislators will continue to develop reactive measures that run the risk of rapid obsolescence as newer technologies are more widely adopted.

MEASURING RISK AND ITS COSTS

Consistent metrics for loss frequency and severity that can communicate the extent of cyber risk are important to understanding how well an organization is prepared for cyber risks. Standardized risk management metrics concepts such as value at risk might also enable organizations to understand the cost and benefits of investing in cybersecurity.

Security always involves tradeoffs (3). The tradeoff might be convenience, money, or technological progress. These tradeoffs can be abstracted into metrics and ultimately used in models to assess cyber risk costs (4). A science of cyber risk will help to address deficiencies in decision support models, which are often formulated ad hoc for individual organizations based on point-in-time data and do not adequately capture cyber risk trends that could influence an organization’s risk management strategy (5). For example, the movement towards cloud providers is a way for organizations to outsource computing security infrastructure. An organization reaps considerable benefits in cost and increased operational efficiency by moving to the cloud. However, if a large number of organizations move to a monoculture of cloud providers, the result could be the exposure of organizations to a new class of systemic risk related to systemic failures across multiple industries. Such risk would not be captured by point-in-time models—instead, analysis requires strategic inter-organization decision science that could be enabled by cross-disciplinary research.

Today, calculating indirect costs of cyber risk such as lost revenues or reputational effects is imprecise at best. There are also gaps in measuring how a cyber event in one department impacts other departments throughout the organization. Current metrics do not capture bigger economic questions such as measuring risk across organizations. Establishing organizational processes for cyber risk planning and enabling these plans with measurable success metrics is another problematic aspect of cyber risk that may also help drive organizational accountability. Non-profit organizations like the Factor Analysis

For Affiliations see Supplementary material.
Corresponding Author: falco@stanford.edu

of Information Risk Institute are working towards developing robust value-at-risk measurements, but more needs to be done as described.

An example of cross-disciplinary collaboration on this topic could be on measuring risk across a portfolio of companies, sectors, or countries. Data scientists could use clustering techniques to determine what other types of catastrophic loss scenarios are most akin to cyber. Then, economists could develop cyber loss scenarios that are consistent with and comparable to the other types of catastrophic event loss scenarios.

CAN CYBER RISKS BE AVOIDED?

Risk avoidance per se is a largely aspirational goal, given the dependence of modern society on digital technology. It is unrealistic to avoid all cyber risk, although it may be possible in some domains to eliminate certain kinds of cyber risk. One option is to design and use inherently secure systems (6). Most devices, networks, and systems, particularly Internet-of-Things (IoT) devices, were built without security as a priority. Some new approaches to designing and building software and hardware systems aim to avoid certain security issues entirely. For example, by designing operating systems with built-in limitations on what processes they can run, certain risks are automatically avoided.

Managerial options to avoid cyber risk include minimizing the use of connected computing systems in certain environments. This is not a decision to be taken lightly, considering the benefits networking affords. However, the thought process behind choosing to unplug is important. Today, society defaults to assuming connectivity is beneficial. Organizations must weigh the tradeoff between convenience and security if they seek to avoid cyber risk.

An example of cross-disciplinary collaboration that would advance understanding of this topic is between political science and management science. Many political scientists have sought to apply nuclear weapons deterrence doctrine to cyber, but the parallels are often weak at best. Deterrence scholars could work with management scientists who have studied non-technical attack/defense dynamics to understand what social dynamics theories may fit with their deterrence conceptualization. Together, they could arrive at alternative mechanisms to avoid risk.

OPPORTUNITIES TO REDUCE RISK

Many cyber security researchers use a framework called the Parkerian hexad to evaluate opportunities to reduce risk. It outlines six elements of information security: confidentiality, possession or control, integrity, authenticity, availability and utility (7). Increasingly, organizations understand that they cannot protect themselves from all risks all the time. They need to prioritize their systems, networks, and data security and evaluate opportunities to reduce cyber risk across the hexad (8). The relative importance of each hexad element must be evaluated.

One concrete step towards reducing cyber risk across the hexad is to share threat information—if an organization shares recent attacks with others that may use similar systems, networks, or data, future attacks against others could be prevented (9). Information Sharing Analysis Centers (ISACs) exist across all critical infrastructure sectors in many countries to facilitate information sharing about cyber events, but they have not been equally successful. ISACs are only as effective as the member organizations' participation. For example the financial services ISAC is well-regarded. One reason for its success is that the banking and financial services regulators in the United States have required participation. Also, U.S. law protects consumers from banking and card services theft, which means that losses from such theft must be absorbed by the industry. This incentivizes companies in financial services to participate. Other ISACs likely will need similar incentives to spur participation. Studying why some sectors engage more effectively with their ISAC, and alternative incentives for participation, could help improve cyber event prevention in other sectors.

In many ways, cybersecurity can be characterized as a public good. This can lead to underinvestment in reducing cyber risk. From the perspective of an individual organization, it might be optimal to reduce investments in cyber risk reduction. In such cases, public policy might be needed to enforce some minimum standards for cyber risk reduction that apply to all organizations in a given sector (10). In principle, these standards need to be global to be effective given that cyber risk knows no boundaries. In light of the lack of policy direction in this space, it is important to evaluate market mechanisms, such as cyber insurance (11) that may have cyber risk reduction effects.

An example of two disciplines that could collaborate here is economists and behavioral scientists. Economists could investigate mechanisms to reduce cyber risk through incentives, which would benefit from working with behavioral scientists providing insight into what drives human interactions and how they engage with technology.

TRANSFERRING RISK

It is important to consider how organizations can transfer cyber responsibility and risk externally to insurers, capital markets, contracts in indemnity and hold harmless agreements, or even the government (12). Cyber risk management today mainly focuses on prevention, while risk transfer instruments like insurance are in their infancy. Insurance coverages are narrow and pricing is based on heuristic estimations of loss frequency and severity. Such risk transfer must be empirically grounded in science in order to accurately price forward-looking cyber risk. Inaccurate pricing for cyber insurance could result in disproportionate payouts compared to premiums paid. This could be especially damaging if cyber risks accumulate, when an insurer faces the possibility of simultaneous payouts on a large scale because of a single event (e.g., an attack on a common technology such as a cloud service provider).

The private sector is unwilling or unable to cover certain cyber events. For example, insurers generally deny coverage for events deemed "cyberwar" due to the potentially large effects of sophisticated nation-state attacks. Scholars need to evaluate government's role for risk transfer in acts of cyberwar. Different organizations and governments have varying perspectives on what is cyberwar, further complicating risk transfer (13). Pooling of cyber risk as launched in Singapore might be a solution to cover risks on a broader scale, involving the private sector and the government. Standardized insurance conditions as used in Germany also help to reduce uncertainty about coverages and make insurance more broadly available.

An example of potential cross-disciplinary collaboration concerning risk transfer could be between law and management science. Legal scholars could investigate contract mechanisms that would enable risk sharing and transfer – research that would benefit from management scientists who could provide insight to how such a contract would impact business operations and a company's balance sheet.

MONITOR AND MANAGE RESIDUAL RISK

Regardless of prevention measures taken, there is still some chance that an organization will experience a cyber event. Monitoring such residual cyber risk often falls to the Chief Information Security Officer (CISO). However, the actual ownership of cyber risk within an organization is generally unclear – even if the organization has a CISO or Chief Information Officer (CIO) who is supposed to manage information security (14). Because cyber risk is a key part of an organization's overall business risk, all departments of an organization have a vested interest in how the organization manages residual cyber risk.

Researchers must devise a systematic approach for business units to increase their involvement in cyber decisions, rather than these responsibilities being centralized only in the CISO or CIO role. Business unit operators need better incident response playbooks, including strategies for communicating with internal units, board members, shareholders, government officials, regulatory authorities, and the public (15).

An example of two disciplines that could work together on monitoring and managing residual risk is computer science and political science. Political scientists are actively studying how cyber capabilities contribute to power dynamics across nations. This would benefit from a computer scientist's ability to actively monitor cyber attacks and capabilities.

CONCLUSION

To overcome barriers to cross-disciplinary cyber risk research, scholars researching a component of this agenda could develop collaborations and explore questions at the margins of the disciplines. An example of where the authors have done this is in studying extreme cyber risk scenarios. Initiated by an economist, we set out to establish a consistent approach for modeling extreme cyber risk scenarios. The team engaged a computer scientist who was able to describe plausible extreme events and the criticality of these events so that the relevant information was analyzed. Neither the economists nor computer scientist had the complete expertise to address the problem.

Another barrier to cross-disciplinary research is the nature of calls for proposals for funding. We cannot assume that cross-disciplinary research projects will all be arrived at organically. Instead, government and private sector research grants for cyber risk must be structured to require collaboration that studies questions on the margins of disciplines. Otherwise, the science of cyber risk will merely borrow from previous studies along disciplinary boundaries – effectively forcing the use of existing methods to imperfectly address cyber issues. For example, a working group at the International Actuarial Association is aiming to develop an economic cyber loss index. To collect incident data, they first need to consult legal scholars to understand what data can be legally collected. Then computer scientists need to collect data using tools at their disposal. Data scientists need to cleanse this data and prepare it for analysis. Management scientists need to categorize the data's relevance by sector, and economists ultimately need to model the economic impact. If any single discipline approached this problem alone, there would likely be missteps throughout that could lead to flawed research outcomes.

Industry and government organizations must demonstrate leadership in setting the direction of cross-disciplinary research. Because digital risk is a daily battle for businesses and public organizations, they must propose and provide data for the most pressing research problems to spur cross-disciplinary scoped projects. One promising concept from the risk management toolbox which is well established in other contexts (natural catastrophe, terror, nuclear risks), might be the pooling of risk. However, to implement this, legal, business, and economic preconditions have to be studied. Further discussion on potential standardization and regulation to improve cyber risk management needs input from various disciplines as well. It is important to understand the economic and political incentives for the proliferation of cyber risk, how society can be trained in dealing with cyber risks, and what legal action can be taken to improve cyber security.

Both the intellectual independence of academia and the experiential knowledge from government and industry are critical for robust, integral research. Integrating practical direction into calls for research proposals that are not catered to a specific discipline can spur cross-disciplinary collaboration, e.g., by agencies such as the U.S. Defense Advanced Research Projects Agency. Encouraging disciplines to pursue their world-view of each question is important, but the disciplines should recognize the limits of their expertise and collaborate towards comprehensively addressing the questions.

REFERENCES AND NOTES:

1. J. J. Cebula, L. R. Young, A taxonomy of operational cyber security risks: Technical note CMU/SEI-2010-TN-028, Software Engineering Institute, Carnegie Mellon University (2010).
2. R. Ramirez, N. Choucri, *IEEE Access* **4.1**, 2216-2243 (2016).
3. G. V. Post, A. Kagan, *Computers & Security* **26.3**, 229-237 (2007).
4. R. Böhme, S. Laube, M. Riek, *Variance Journal* **12.2**, 161-185 (2019).
5. L. A. Gordon, M. P. Loeb, The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* **5.4**, 438-457 (2002).
6. H. Shrobe, D. L. Shrier, A. Pentland, eds. *New Solutions for Cybersecurity*. MIT Press (2018).
7. A. Beaument, D. J. Pym, *WEIS*, 1-20 (2010).
8. G. Falco, C. Caldera, H. Shrobe, *IEEE Internet of Things Journal* **5.6**, 4486-4495 (2018).
9. S. Laube, R. Böhme, *ACM Computing Surveys (CSUR)*, **50.5**, 77:1-36 (2017).
10. D. Thaw, *Ga. St. UL Rev.* **30.2**, 287-374 (2014).
11. M. Eling, W. Schnell, *Journal of Risk Finance* **17.5**, 474-491 (2016).
12. L. Bodin, L. A. Gordon, M. P. Loeb, A. Wang, *Journal of Accounting and Public Policy* **37.6**, 527-544 (2018).
13. D. Woods, A. Simpson, *Journal of Cyber Policy* **2.2**, 209-226 (2017).
14. V. Hooper, J. McKissack, *Business Horizons* **59.6**, 585-591 (2016).
15. G. Falco, A. Noriega, L. Susskind, *Journal of Cyber Policy* **4.1**, 90-116 (2019).

Acknowledgements: See supplementary materials for author disclosures.