

A Fundamental Approach to Cyber Risk Analysis

Rainer Böhme, Stefan Laube, Markus Riek

Universität Innsbruck, Austria

Corresponding author: rainer.boehme@uibk.ac.at

Abstract. This article provides a framework for actuaries to think about cyber risk. We propose a differentiated view on cyber versus conventional risk by separating the nature of risk arrival from the target exposed to risk. Our review synthesizes the literature on cyber risk analysis from various disciplines, including computer and network engineering, economics, and actuarial sciences. As a result, we identify possible ways forward to improve rigorous modeling of cyber risk, including its driving factors. This is a prerequisite for establishing a deep and stable market for cyber risk insurance.

Keywords. Cyber risk, information technology, networks, economic modeling

Contents

1	Introduction	3
2	Cyber Risk	4
2.1	Technological Context	4
2.2	Defining Cyber Risk	8
2.3	Insuring Cyber Risk	10
3	Cyber Risk Management	13
3.1	Risk Factors	13
3.1.1	Threats	13
3.1.2	Vulnerabilities	15
3.1.3	Controls	16
3.1.4	Assets	17
3.1.5	Impacts	17
3.2	Determinants of Risk Factors	18
3.3	Cyber Risk Management in Practice	19
4	Economic Modeling	21
4.1	Framework	21
4.2	Common Assumptions	23
4.3	Information Asymmetries	26
4.3.1	Adverse Selection	26
4.3.2	Moral Hazard	27
4.3.3	Insurance Fraud	28
4.4	Interdependent Security	28
4.4.1	Model and Example	29
4.4.2	Implications	30
4.4.3	Literature	31
4.5	Market Structure	32
5	Actuarial Modeling	33
5.1	Data Sources	33
5.2	Individual Loss Distribution	35
5.3	Compound Loss Distribution	38
6	Conclusion	41

1 Introduction

This article sets out to provide a framework for actuaries to think about cyber risk. With the ambition of being fundamental, the article tries to establish a distinct notion of “cyber” which can be associated with a set of characteristics relevant to conceptual as well as quantitative modeling of cyber risk. To this end, the article identifies and explains important factors that affect loss distributions after cyber attacks and, by extension, market participants’ decisions to offer or seek insurance for cyber risk. It reviews selected scholarly works that apply economic and actuarial concepts to the domain of cyber risk.

Understanding cyber risk is a hard problem. Standard textbooks on technical aspects of information security have many hundred pages in dense technical jargon and still cover each topic only superficially. Even seemingly simple security mechanisms used by billions of people everyday are not fully understood. For example, according to Bonneau et al. (2012), a few hundred research papers address the problem of password security. The number of data points analyzed in these works accumulates to many millions. Yet, this issue is far from being solved. Adding a risk perspective potentiates the problem space by including the combinatorics of technical, social and economic factors. This article provides a first step in structuring the field, by making it accessible to analysts trained in more conventional domains of risk modeling and highlighting the specifics of cyber risk.

Our approach is as follows. We first develop a definition of cyber risk useful for the insurance industry (Section 2). This involves recalling the key concepts of information technology and how they shape risk. Then we discuss the treatment of cyber risk from three relevant angles: cyber risk management by firms and organizations (Section 3), economic modeling of cyber risk transfer (Section 4), as well as actuarial modeling of cyber risk (Section 5). We close by discussing practical challenges to cyber risk insurance and promising research directions (Section 6).

The target audience of this article are trained risk analysts and actuaries in the insurance industry. Some aspects may also be relevant for underwriters, brokers, corporate risk and information security managers, policy analysts, as well as academics studying related fields. The article addresses topics with global scope. Therefore we refrain from making policy recommendations which would be specific to a few jurisdictions.

2 Cyber Risk

Undoubtedly, the term *cyber* has become a buzzword among marketers and policy makers, causing a plethora of associations, but lacking a single unanimous definition. The application of the term cyber is vast, from the “cyber space” popularized after the science fiction short story “Burning chrome” (Gibson, 1987) to the establishment of the “US Cyber Command” in 2010. The latter is a consequence of recognizing cyber to be the fifth domain of warfare next to land, sea, air, and space (Lynn, 2010). What matters from a risk and insurance perspective is whether the qualifier “cyber”, for instance in cyber crime, is merely an exchangeable technology prefix, such as “electronic” (or shorthand “e”) in electronic commerce; or if it can meaningfully demarcate cyber risk as a new class of risk that requires special treatment and tailored conceptual models.

2.1 Technological Context

In order to distinguish cyber from non-cyber (or *conventional* as a qualifier for risk and insurance), it is helpful to recall key characteristics of the enabling technology. All advances in information technology result from progress in microchip manufacturing, chiefly the technical capability to densely integrate electrical *circuits* in mass production, along with the engineering tools to design circuits with millions of logical elements (gates) in a systematic and largely predictable manner. The design complexity of circuits hints to the

first source of risk. The number of possible states of a few hundred gates exceeds the number of atoms in the known universe, and the relation between inputs, states, and outputs of circuits can be highly non-linear. Therefore it is impractical to test and fully predict the behavior of arbitrary circuits built with current technology even if their design is known.¹

Microchip production is characterized by substantial economies of scale. Creating the tool (a set of masks) to produce the first batch of a circuit is very expensive; the cost for every following batch is negligible compared to the upfront cost. The industry has adapted to this cost structure. It mainly produces general purpose circuits, *programmable* to control many different applications. This highlights the crucial role of software in determining the behavior of most electronic devices. Take for instance electronic elevator control units. In the 1970s, these devices were built solely for this purpose and designed with exactly the components needed to realize the required logic. The resulting circuits were complicated, but it was possible for a trained expert to understand the system completely. The design complexity still allowed to analyze error sources in case of incidents and possibly infer about causes and effects. Understanding these relations is essential for risk analysis and the attribution of losses. Modern control units are built from general purpose microprocessors, which densely integrate zillions of electronic components. The behavior of such systems is defined by customized *software*, which implements the required logic and often only uses a small fraction of the functionality offered by the general purpose component. The resulting systems exhibit high design complexity. This increases the risk that design flaws remain unnoticed and raises the effort of proactive risk research as well as the cost of forensic analyses after incidents. Both add to uncertainty and raise the transaction costs of risk transfer arrangements.

These are not the only issues software systems cause from an insurance perspective. It

¹ It is possible to design circuits and software with automated verification in mind. This technology is orders of magnitude more expensive than common personal computer (PC) or smartphone technology and thus exclusively used for critical applications, such as passenger aircrafts or nuclear power plants.

is easy to revise software, as compared to upgrading hardware (Shapiro and Varian, 1998). This leads to short product life cycles and consequently less time to collect actuarial data for a system in operation. Possibly most worrying is the fact that some classes of common programming errors allow malicious third parties to reprogram the logic of a device. As many devices are over-equipped for their purpose, the existence of hidden functionality is often hard to detect because the modified device still behaves as expected. We have witnessed such hostile take-overs in the form of computer viruses on personal computers and malware infections on smartphones. The Stuxnet worm (Chen, 2010), as an early example of state-sponsored cyber warfare, highlights that programmable devices in industrial control systems are not exempt. With the proliferation of programmable devices in every aspect of life – typically referred to as the Internet of Things (IoT) – it is only a matter of time until cyber criminals enter these platforms.

The ongoing connection of programmable devices to *networks* with wired or wireless physical links leverages the risk to be considered by insurers mainly through two channels. First, compound systems consisting of many interconnected microprocessors, often owned and controlled by different parties, exhibit additional design complexity. In many cases no entity has a global plan or view of the overall system, rendering validation and prediction very hard or impossible. Second, networking increases the surface for and reach of malicious attacks. While modifying the software of standalone devices requires physical proximity at least once in the lifecycle (including the supply chain), networked devices, if not sufficiently secured, can be reprogrammed remotely from any other device in the network. Consequently, devices connected to the Internet, an inter-network designed for routing data packets globally, are in principle exposed to threats from any other person with access to the Internet.

Networked systems further complicate risk analysis because insurers are rarely in the position of insuring a network as a whole. Instead, different parts of the network (nodes)

are operated by autonomous decision makers (agents), each with different interests (utility functions), information sets (information asymmetries) and expectations. They make independent economic decisions while being connected by a common factor that affects the joint outcome (externalities). Part of the agents' decision space is whether to invest in security (risk mitigation) or buy insurance (risk transfer).² The economic terms in parentheses suggest that cyber risk analysis is as much about understanding and modeling the technology as it is about understanding and modeling the economic incentives of the involved agents (Anderson and Moore, 2006).

But technological advances also yield new opportunities for insurers. This generation is about to witness a digital revolution because the very same technology opens a large space for innovation, often with the potential of unleashing unprecedented economic growth (e. g., Brynjolfsson and Hitt, 2003). To continue the previous example, the microprocessor powering a digital elevator control has spare capacity to solve more complicated logic. For instance, it could be programmed to predict demand or to coordinate with a second elevator in the same building; functions that are typically marketed with the keyword “smart”. Other examples of recent innovations, such as autonomous vehicles, would not be possible without microprocessors and networks that supply relevant data on request, e. g., traffic information. As a result of these and similar developments across many industrial sectors, businesses, households, and governments increasingly depend on information technology in numerous ways (Brynjolfsson et al., 2002). What is a center of value creation turns into a loss center at the moment when the technology fails to serve its purpose. This allows insurers to offer coverage for technology and assets that depend on it.

² The canonical risk management instruments further include risk avoidance and risk acceptance.

2.2 Defining Cyber Risk

We have used the term “risk” informally in the previous section. For our definition of cyber risk, it is convenient to start with a high-level notion of conventional risk inspired by the management literature (e.g. Kaplan and Garrick, 1981):

$$\text{Risk} = \text{Probability of a loss event} \times \text{Magnitude of loss} . \quad (1)$$

This definition of risk has many deficits. It coerces a complicated loss distribution to a single Bernoulli trial, is agnostic about the time dimension, and does not differentiate between individual and aggregated losses caused by a single loss event. Yet, this definition is useful to motivate a differentiation that helps to develop a more precise notion of cyber risk than commonly used. We propose to distinguish between

- *risk arrival*, the processes causing loss events, modeled in the simplest possible form as probability of a loss event in Equation (1), and
- *target*, the assets that suffer losses, modeled as a fixed magnitude of loss in Equation (1).

Both risk arrival and target can independently belong to one of the two classes, *cyber* or *conventional*. We classify risk arrival as cyber if the loss event is primarily caused by logic (and by extensions: computer programs and networks), and as conventional if the loss event is primarily caused by physical force. Likewise, we speak of a cyber target if the loss event predominantly devalues information assets, e.g., through destruction (data loss or unauthorized modification), loss of exclusivity (data breach), and consequences thereof (data abuse). A conventional target incurs losses by the destruction or unavailability of physical assets. Borderline cases exist along both dimensions. However, we believe that this distinction is useful, for instance because insurers can implement a modular approach

with division of labor: one team specializes on the specifics of cyber at the risk arrival process; another team of experts values cyber assets and quantifies losses. With this distinction, we can precisely classify all cyber incidents proposed in Box 1.

Box 1: Classification of cyber incidents

Example 1 *Ransomware encrypts clients and servers used by the software development team. All work carried out since the last backup is lost.*

In this example we have cyber risk arrival and a cyber target.

Example 2 *An earthquake destroys a data center.*

In this example we have conventional risk arrival and a cyber target.

Example 3 *A distributed denial-of-service (DDoS) attack against a major airport interrupts and delays air traffic in a geographic region.*

In this example we have cyber risk arrival and a conventional target.

We have collected alternative definitions of cyber risk to cross-check our definition. Most authors agree that the involvement of networked computers is an essential element (cf. Ögüt et al., 2011; Böhme and Schwartz, 2010). Among the more specific definitions, we observe that some emphasize the cyber element in the arrival process (cf. Anderson et al., 2008; Mukhopadhyay et al., 2013; Stoneburner et al., 2002) whereas others focus on cyber targets (cf. Biener et al., 2015; Cebula et al., 2010; ISO/IEC, 2014; Eling and Schnell, 2016). This observation confirms our belief that a unified definition should include both aspects and weigh them equally.

Our approach is compatible with conventions in the literature to classify cyber crimes by the role of computers and networks in criminal acts (cf. Goodman, 1997; Alkaabi et al., 2011). Computers and networks can be targets, facilitating tools, or incidental aspects of crimes. The third category is dropped in recent definitions of cyber crime since the ubiquity of the Internet makes computer networks an incidental aspect of almost any crime.

We have also considered alternative distinctions, such as the difference between tangible and intangible losses or the difference between random failures and malicious attacks. As these attributes are known for conventional insurance and do not characterize the specifics of cyber risk, we consider them as conceptually orthogonal to our distinction by the domain of risk arrival and target.

2.3 Insuring Cyber Risk

Cyber insurance is a vehicle for cyber risk transfer. In exchange for a defined premium and for a defined period of time, the insurer contractually agrees to financially compensate potential losses incurred by the insured through the realization of cyber risk. Like in conventional insurance, losses may comprise primary and secondary losses (the sum of which is referred to as recovery cost) and the indemnity may include first party as well as third party losses for which the insured is held liable (Anderson et al., 2008, p. 82). Box 2 summarizes the state of the cyber insurance market at the time of writing.

Box 2: State of the cyber insurance market

The size of the global cyber insurance market is hard to estimate as insurers do not share detailed information on premiums and claims. Some industry reports give an impression of the state of the market in 2015 and 2016:

- AGCS (2015): Gross written premiums estimated at \$ 2 billion;
- Betterley (2016): Gross written premiums estimated at \$ 3.25 billion (US);
- NetDiligence (2015): Total claim payouts estimated at \$ 1.5 billion.

According to AGCS (2015), the US market accounts for the majority of the gross written premiums. The European market is expected to catch up, driven by new breach notification laws that may incentivize incident reporting. To put the premiums into context: NetDiligence (2015) reports a total claim payout of \$ 75.5 million based on 132 claims, which constitute approximately 5% of all claims.

Using the definition of cyber risk from Section 2.2, we differentiate between three forms of

		Target	
		Physical assets	Information assets
Risk arrival Logic Force		Conventional insurance	Cyber-asset insurance
		Cyber-threat insurance	Cyber insurance

Figure 1: A classification of cyber risk insurance

cyber insurance as illustrated in Figure 1. The figure tabulates the domains of risk arrival in rows and the domains of the target in columns. Conventional (non-life) insurance is located in the top left area. Cyber insurance in a narrow sense is located in the bottom right area. This is what most authors of early contemplations of cyber insurance presumably had in mind (e. g., Medvinsky et al., 1994; Schneier, 2001; Grzebiela, 2002; Baer, 2003).

Several reasons explain why the market for cyber insurance in a narrow sense did not evolve as predicted, chiefly a lack of demand and a lack of claims in the 1990s. The lack of claims was interpreted as an indication of cumulated risk. This spurred fears of a “cyber hurricane” which led reinsurers to stop covering cyber risk in the early 2000s (Böhme and Schwartz, 2010). While this market did not evolve as predicted, nascent markets for cyber-threat and cyber-asset insurance exist.

The increasing demand for cyber-asset insurance, in the top right corner of Figure 1, is a consequence of the growing dependence of organizations on information systems and the data processed therein. While some industry experts initially saw difficulties in writing precise policies for the intangible losses of information assets, the progress made in valuating intangible assets in the finance and accounting disciplines has accommodated this concern. A major development to this end is the uptake of marketplaces for IT services, for example in various cloud computing models, and for business data (Balazinska et al., 2011; WEF, 2011). Both market types generate price information useful for estimating the monetary

value of (lost) in-house infrastructure or databases.

Another avenue for the insurance industry to get in touch with cyber risk is the bottom left corner of Figure 1, provisionally termed cyber-threat insurance. Although cyber attacks are commonly excluded from property or liability insurance, insurers and reinsurers begin to realize that they might be exposed to cyber threats indirectly through business interruption policies. Networked information technology has become so vital, in particular for operations and the supply chain, that computer networks are an incidental aspect of almost every business interruption. Policies with explicit exclusions of cyber will therefore become harder to sell. We observe a trend that exclusions will be defined more narrowly, for instance specifically naming acts of cyber warfare or cyber terrorism, and combined with a tight limit in the order of \$50 million for incidental cyber risk not belonging to the specified categories. This also limits the exposure if an act of cyber warfare cannot be attributed unanimously to a nation state (Clarke and Knake, 2012), a situation experienced with the Sony hack in 2014.³

Another line of business is cyber insurance of third party risk. In that case policy holders are held liable for monetary amounts which can be passed on to the insurer. This partly explains the success of insurance policies covering losses from breach reporting obligations (Romanosky, 2013; Kirkpatrick, 2015; Laube and Böhme, 2016; Bandyopadhyay et al., 2009). Corresponding legislation is high on the policy agenda around the globe in order to reduce under-reporting of cyber attacks. However, important limitations remain. Most policies cover only parts of the costs, define rather tight limits on each loss source, and exclude losses of reputation and subsequent negative stock market reactions, which are hardly measurable (Acquisti et al., 2006). Moreover, insuring against regulatory fines is prohibited in many jurisdictions (Gatzlaff, 2012). If quantifying losses ex-post is diffi-

³For further details on the hack at Sony Pictures, see <https://www.washingtonpost.com/news/the-switch/wp/2014/12/18/the-sony-pictures-hack-explained/>.

cult, insurer and policyholder may agree on a fixed indemnity ex-ante, provided that the agreement complies with the codification of the indemnity principle in the jurisdictions concerned.⁴

3 Cyber Risk Management

In order to insure cyber risks, insurers must be able to identify essential risk factors and understand the decision of firms to seek insurance for specific risks.⁵ To provide a fundamental background, this section describes risk factors using the typical chain of causality from cyber threats to financial losses, and recollects selected aspects of how firms implement information security in practice.

3.1 Risk Factors

We conceptualize a cascade model of cyber risk arrival, depicted in Figure 2, to organize five classes of risk factors: threats, vulnerabilities, controls, assets, and impacts. Our model is tailored to the scope of this article. For a more comprehensive path model see Ransbotham and Mitra (2009).

3.1.1 Threats

Threats subsume accidental physical or logical errors and intentional action by malicious attackers. Our model defines threats as the root causes of loss events, although it is possible in principle to dig deeper. As cyber crime became a profit-driven industry in the early

⁴ The principle of indemnity for non-life insurance states that insureds must not profit from the occurrence of a loss event (cf. Mehr and Cammack, 1972). Hence, an upper bound for the compensation is the actual loss incurred.

⁵ On terminology: we use *firms* as shorthand for professionally managed organizations in the private or public sector exposed to cyber risk. Firms appear on the demand side of insurance markets. *Policyholders* are firms who have acquired coverage from *insurers*. *Vendors* supply information technology and software to firms.

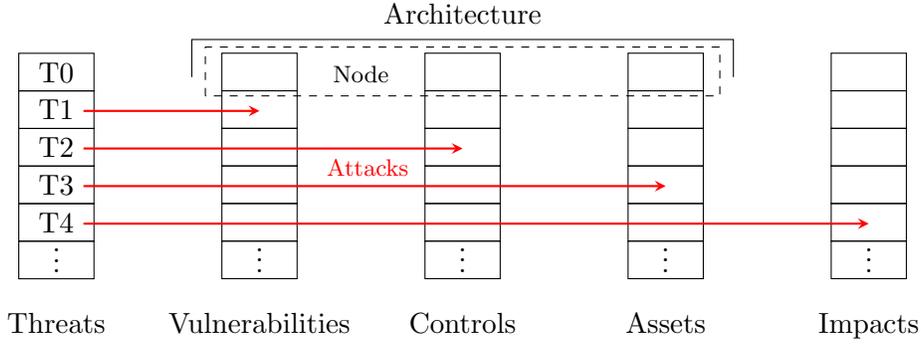


Figure 2: Cascade model of cyber risk arrival

2000s (Anderson et al., 2008), scholars have started to analyze the motivations of attackers drawing on criminology (Cressey, 1953) or the economics of crime (Becker, 1974). Attackers use various methods including technical means, such as malformed network packets, non-technical scams, such as social engineering to obtain access credentials, or a combination of both.

The conservative assumption that attackers always choose the *weakest link* has almost become a mantra in security trainings. A more precise model is to broadly classify attack strategies by their relation to victims (Herley, 2014). *Opportunistic* attackers do not care about whom they attack. They standardize attack techniques to enjoy economies of scale when hitting several targets, of whom only a small fraction gets victimized. Unsolicited mass email (spam) is an example for this strategy. By contrast, *targeted* attacks focus on a specific victim and customize the attack method. They often involve several iterations of information gathering, with the aim to maximize the success probability against a carefully selected target. This strategy pays off for high-value targets, such as in extortion and industrial espionage. An important implication for risk analysis is that modeling threats as probabilistic is more appropriate for opportunistic attackers than for targeted attacks. More adequate analytical tools for the latter are worst case approximations and game

theory.

3.1.2 Vulnerabilities

Not every threat realizes a risk. Threats require *vulnerabilities* in the target system to become successful attacks. To continue the examples from above, the malformed network packet is only harmful if the software processing the data packet enters an undefined state which allows the attacker to take over control.⁶ Such vulnerabilities emerge from common programming mistakes, which are hard to be fully avoided in the software development process. Likewise, the social engineering attempt is only successful if the victim is tricked into sharing credentials with unauthorized parties. With respect to Bandyopadhyay et al. (2009), we may distinguish between *symptomatic* and *systemic* vulnerabilities. The former affect only a single firm (e. g., because it uses custom software or runs an erroneous configuration) whereas the latter exists in many firms (e. g., when standard software is vulnerable or a default password has not been changed).

Technical vulnerabilities matter for risk analysis in two ways. First, systemic vulnerabilities expose many targets to the same threat. With the ability to scale attacks by automating them on programmable devices and using networks as propagation vectors (cf. Section 2.1), many firms are at risk of suffering losses at the same time. This correlation between risks leads to fatter tails of the cumulated loss distribution and may hamper insurability (Böhme and Kataria, 2006). Second, information about vulnerabilities is notoriously incomplete, leading to a race for information between attackers and defenders (Ransbotham et al., 2012). There is some controversy about the right regime of distributing vulnerability information between stakeholders. Some consider vulnerabilities a strategic asset for national security (Moore et al., 2010), others a tradable information good (Böhme, 2006). The distribution regime impacts the vulnerability discovery process

⁶ In most cases, other parts of the data packet contain instructions that reprogram the software.

by setting incentives for security researchers. It may also affect vendors' efforts to produce software with fewer vulnerabilities and distribute patches. In practice, many disclosure regimes coexist, with underground markets on one end of the spectrum and organizations committed to responsible disclosure on the other end (Miller, 2007; Arora et al., 2008; Zhao et al., 2015). Given the decisive role of vulnerability information, the insurance industry will likely take part in the vulnerability ecosystem when it covers substantial amounts of cyber risk.

3.1.3 Controls

Not every pair of threat and vulnerability leads to a successful attack. Firms can place technical and non-technical *controls* to mitigate cyber risks. Awareness campaigns and trainings for employees are examples for non-technical means. Technical controls can be *detective* controls, which indicate the realization of threats, possibly trigger alarms, and require reaction, or *preventive* controls, which proactively shield specific vulnerabilities from threats (cf. Cavusoglu et al., 2004). Regarding preventive controls, most people think of add-ons, such as network packet filters or antivirus software. However, structural changes, such as dedicated networks for critical data disconnected by “air gaps”, are preventive controls, too. All controls cause costs which firms must weigh against the expected benefit in terms of prevented losses. Many scholars have studied this decision problem through the lens of investment theory (cf. Hoo, 2002; Gordon and Loeb, 2002; Su, 2006; Böhme, 2010). The management of controls often requires strategic decisions, such as prescribing minimum access rights to information or specifying a patch strategy. In the latter case firms face a trade-off between fast roll-out of software updates to close known vulnerabilities and the risk that barely tested patches have unanticipated side-effects or break critical business processes (Beattie et al., 2002; Ioannidis et al., 2012).

All types of controls are relevant for risk analysis. Detective controls improve informa-

tion and provide a more direct access to data on the risk arrival process than counting (aggregated) losses. Preventive controls mitigate the risk of specific threats.

3.1.4 Assets

The interaction between threats, vulnerabilities, and controls determines the success of attacks. Attacks turn into incidents if they hit critical *assets*. It is not necessary that the asset is damaged or destroyed; a customer database leaked to outsiders may be as painful as losing it entirely. In general, any undesired breach of a canonical protection goal (confidentiality, integrity, availability) is considered a security incident. Including assets as an individual risk factor is important because firms may use similar technology to secure assets of very different value and sensitivity. To illustrate this, recall that the technology behind secure Internet connections, the HTTPS protocol, does not substantially differ between online banks and well-administered websites of online pizza shops.⁷

For actuaries the valuation of assets (as discussed in Section 2.3 in the context of *cyber-threat insurance*) is an essential prerequisite to the estimation of expected losses if any of the different protection goals is violated.

3.1.5 Impacts

The value and criticality of affected assets influences the potential *impact* of an incident and thus the amount claimed under an insurance policy. The impact may exceed the asset value by orders of magnitude, for instance if a firm stores customers' passwords in plaintext⁸ and faces substantial liability and compliance cost along with severe reputation damage after a breach exposed customer data to the public. The impact can also be just a fraction of the asset's value, if an effective incident response and recovery prevent larger

⁷ One may still hope that the bank has better processes and hires more qualified staff to look after its security.

⁸ There is no benefit in storing passwords in plaintext. Best practice is to transform them with specialized cryptographic one-way functions including a different constant ("salt") for each user account.

damage. Insurers may provide professional incident response services for their policyholders to reduce impacts (AGCS, 2015).

3.2 Determinants of Risk Factors

We have introduced the cascade model in order to structure the factors to be considered when modeling cyber risk arrival and loss distributions for different targets. The model also allows a risk analyst to reason on what drives these factors. Most experts will accept threats as exogenous or environmental factors although, in principle, policy or industry initiatives can try to tame the threat environment (Asghari et al., 2015). In general, threats follow global trends in the long run, and attacker tactics in the short run.

By contrast, vulnerabilities, controls, and assets are predominantly endogenous factors, i. e., firms can in principle control them. These three factors together constitute a firm's information *architecture* (as annotated in Figure 2). However, not every firm can exercise direct control on all three factors. Many vulnerabilities originate in software or components supplied by external vendors. With high concentration in software markets and the aim for standardization and interoperability, many firms have little choice about this part of their architecture (cf. Carr, 2003). Similarly, the type and value of assets largely depends on the firm's business, size, and level of technology adoption. By contrast, controls are largely in the responsibility of individual firms. Therefore risk analysis must consider security investments as a relevant determinant of cyber risk. Because firms decide strategically on their security investment, insurers need to understand firms' incentives in order to prevent adverse selection and moral hazard (see Sections 4.3.1 and 4.3.2 below).

Firms cannot always control the impacts of a successful attack. While plaintext passwords in the example above could have been avoided with the right controls, it is easy to find examples where the impact is driven exogenously. A breach may lead to more adverse public reaction if a firm is the first or the only one affected in an industry. Post-breach

crisis communication and incident response may matter a lot. For example, an industry report estimates the cost of data breaches per record between a few cents and \$1.6 million (NetDiligence, 2016). The orders of magnitude difference hints to the imponderables with this type of risk. A particular challenge for modeling the impact of privacy breaches is that sensitive information often results from joining leaked records with other public or proprietary databases available at the time of the breach or in the future (Sweeney, 2002; Narayanan and Shmatikov, 2008). The likelihood and consequences of such events are very hard to predict.

3.3 Cyber Risk Management in Practice

From an insurance perspective it is important to understand how (well) cyber risks are managed in firms and where to find such information. Firms synthesize the canonical instruments from general risk management and IT security to conduct cyber risk management. Cyber risk management requires the definition of a security policy, which guides a firm's information security operations. This, in turn, is centered around the problem of managing and enforcing authorization decisions in a systematic and efficient manner.

Approaches to embed risk management in the organizational structure of firms range from a single centralized department to delegated officers in each business unit. Intermediate forms also exist. The key roles that enable cyber risk management are (1) the senior management, approving an organization's security policy, (2) the chief information security officer (CISO), as the person in charge of a firm's cyber security, and (3) security specialists (often embedded in IT operations), carrying out security-related tasks (Stoneburner et al., 2002). Issues may arise if the CISO does not have direct access to the senior management, or if a source of cyber risk is located outside of the officer's area of influence.

Several standards support risk management in organizations. They can roughly be divided into two classes: top down (focusing on the risk management process) and bottom

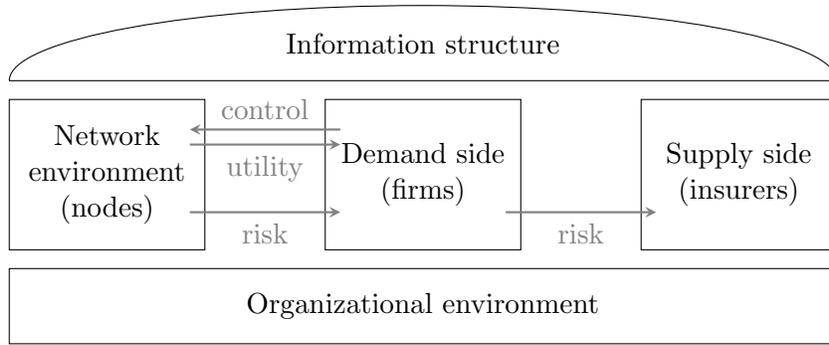


Figure 3: Framework for economic models of markets for risk cyber insurance (Böhme and Schwartz, 2010)

up (focusing on identification and quantification of individual risks). A prominent example of the former category is the ISO 27000 series (ISO/IEC, 2014), adopted by many large firms across different sectors (PwC, 2016). The OCTAVE methodology is an example for the latter (Caralli et al., 2007). Moreover, industry-specific standards, such as the Payment Card Industry Data Security Standard (PCI-DSS), certify compliance with defined practices (PCI, 2015). From an insurance perspective, standards and certifications of firms are indicators of good security practices. However, industry experts complain that many standards and certifications are incoherent, and compliance with them is a weak predictor of actual security. For instance, the retailer Target was certified to be PCI compliant just before being hit by a major breach.⁹ Furthermore, voluntary certifications provide weak signals because they may suffer from adverse selection: less secure firms have stronger incentives to seek certification (Edelman, 2011).

⁹ For further information, see <http://www.darkreading.com/risk/compliance/target-pci-auditor-trustwave-sued-by-banks/d/d-id/1127936>.

4 Economic Modeling

Most economic models of cyber risk insurance analyze if and under what conditions markets for cyber risk insurance can exist. This depends on the properties of the loss distribution and the actions of all market participants involved. We adapt the framework of Böhme and Schwartz (2010), depicted in Figure 3, to structure the action space of all involved parties.

4.1 Framework

The framework has two natural components, which correspond to the demand and supply side of insurance markets. The specific characteristics of cyber risk are modeled by the component called *network environment*. The network environment is composed of atomic but connected elements called *nodes*. Generally, firms control nodes to extract utility from the network. As a side effect, they are exposed to risk that arrives as described by the cascade model of Section 3.1. It is convenient to assume that each node represents exactly one cascade of risk factors (as annotated in Figure 2). Firms control many nodes, interconnected by physical, logical, and social links, such as trust relationships. The links between nodes are typically represented as edges of graph structures.

An obvious example of cyber risk dependencies in a network environment is the propagation of malware, which infects several networked computers (nodes). However, risks can also arise in the software engineering process, as the security of a software product depends on all individual components. From this perspective connected nodes can model single points of failure due to vulnerabilities in common components or dependence on a single external supplier, for instance a cloud service provider. Risks can also propagate in systems of an interconnected supply chain or, from the social perspective, by sharing confidential information with business partners.

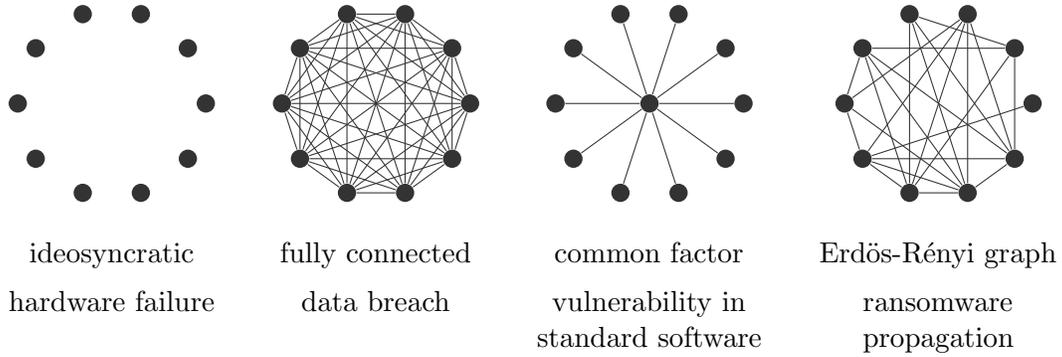


Figure 4: Topologies of connected nodes modeling the dependence in cyber risk arrival for selected threats

This list of dependence relationships is clearly not exhaustive and each type of relationship is characterized by a specific graph topology (see Figure 4). Conceptually, every real-world insurance policy bundles the risk exposure of many nodes; typically those controlled by the policyholder, or a subset if exclusions apply. These nodes possibly maintain connections to nodes under the control of other firms. This assumption naturally generalizes the model to a setting with interdependent security, a specific type of externality discussed later in Section 4.4.

The known issues of insurance markets related to information asymmetries, in particular adverse selection and moral hazard, can be incorporated by instantiating another component of the framework called *information structure*. Information on cyber risk is inherently asymmetric. The complexity alluded to in Section 2.1 precludes even the vendors of programmable components from fully predicting their behavior when combined in practical systems. The information structure bundles modeling decisions that affect the distribution of information (including remaining uncertainty) about the state of the world among the economic agents.

The last component, *organizational environment*, defines general rules. It comprises all parties who can intervene with the cyber insurance market although they do not directly

appear on the demand or supply side. In a model world where market participants follow rational choice, possible interventions take the form of adjusting incentive structures such that a desirable social outcome becomes more likely. An example is regulators' introduction of breach notification laws to incentivize security investments at firms.

4.2 Common Assumptions

Most economic models emphasize the risk arrival process in a single time period short enough to neglect discounting. A common assumption is to model the loss as binary outcome: either the firm faces a loss of fixed size $l > 0$ or no loss at all. The loss amount l is measured on a monetary scale. This reduces the target dimension of cyber risk to a Bernoulli random variable $L \sim B(1, p) \cdot l$, where $B(n, p)$ is the binomial distribution parameterized by the number of controlled nodes n and the probability of loss p . Values of $n > 1$ are reserved for situations where firms control multiple nodes, often in combination with variants of the binomial distribution where trials are not independent.

With $n = 1$ and l fixed,¹⁰ the only way to model control over the risk exposure is to let a firm's probability of loss p depend on its risk mitigation efforts. This leads to a monotonically decreasing *defense function* mapping a monetary *security investment* s to a probability of loss p .¹¹ A simple defense function is given by

$$p = D(s) = \beta^{-s} , \tag{2}$$

where $\beta > 0$ is a parameter of *security productivity*. It controls how efficient money spent on security controls can prevent threats from materializing as successful attacks. More complicated form of defense functions can be found in the seminal literature on

¹⁰ Some authors normalize all monetary values to the unit of the loss, allowing them to define $l = 1$.

¹¹ Security investment corresponds to the concept of self-protection in the sense of Ehrlich and Becker (1972).

information security investment (Gordon and Loeb, 2002). Most defense functions are convex, suggesting an interpretation of decreasing marginal benefit of security investment.

The rational choice paradigm predicts that firms maximize their expected wealth by choosing the security investment level $s^* \leq l$ which minimizes the sum of expected losses and the associated cost of controls s :

$$s^* = \arg \min_s (D(s) \cdot l + s) . \quad (3)$$

This equation models the optimal decision for *risk neutral* firms. Insurance markets can only exist if firms are *risk averse*, meaning that they are willing to trade expected wealth for reduced variance. Economists model risk aversion with concave utility functions¹² $U_\sigma : \mathbb{R} \rightarrow \mathbb{R}$, which map monetary wealth to preference scores (Pratt, 1964). Parameter $\sigma > 0$ controls the strength of the risk aversion. Firms' preferences are ranked by expected utility – which differs from the utility of the expected wealth – and the utility has to be computed over the total wealth including security investment. Risk averse firms optimize the following problem:

$$s_\sigma^* = \arg \max_s (E [U_\sigma(W_s)]) , \quad (4)$$

where W_s is a random variable of the firm's final wealth given security investment s , and $E[\cdot]$ is the expected value operator. For the simple Bernoulli loss model and initial wealth w , we obtain

$$s_\sigma^* = \arg \max_s \left(D(s) \cdot \underbrace{U_\sigma(w - l - s)}_{\text{bad case: loss}} + (1 - D(s)) \cdot \underbrace{U_\sigma(w - s)}_{\text{good case}} \right) . \quad (5)$$

¹² Classes of utility function vary in whether and how the preference for certainty depends on initial wealth. Utility functions with the constant relative risk aversion (CRRA) property are independent of a firm's initial wealth.

Market insurance can be modeled by allowing firms to buy insurance for losses up to a limit x at the price of a linear premium $\pi \cdot x$. In the Bernoulli loss model and respecting the indemnity principle, it must hold that $x \leq l$. With optional market insurance, firms optimize their action in the (s, x) plane:¹³

$$(s, x)_\sigma^* = \arg \max_{s, x} (D(s) \cdot U_\sigma(w - l - s + (1 - \pi) \cdot x) + (1 - D(s)) \cdot U_\sigma(w - s - \pi \cdot x)) . \quad (6)$$

Insurers can only offer this policy if they have complete information and all firms are homogenous. In this case, insurers anticipate firms' choices and adjust π to reflect the risk given the common security investment s^* :

$$\pi = D(s^*) \cdot (1 + \lambda) , \quad (7)$$

where λ is a *loading* needed to pay administrative costs, the cost of safety capital, and insurer's profit. A premium is called *actuarially fair* if $\lambda = 0$, which implies full insurance $(s, x)_\sigma^* = (0, l)$.¹⁴ It also highlights that insurance and security investment are substitutes in this model. In practice, actuarially fair premiums are only sustainable with subsidies.¹⁵ If $\lambda > 0$, the existence of an insurance market (i. e., $x_\sigma^* > 0$) depends on the strength of risk

¹³ Even simpler models producing the same qualitative results restrict one or both choice variables to two options, $s \in \{0, \mathfrak{s}\}$, $x \in \{0, l\}$. This simplifies the optimization problem to a comparison of cases.

¹⁴ Because firms can eliminate all variance without sacrificing expected wealth. Marginally risk averse decision makers prefer this situation. Security investment must be zero because otherwise firms could improve their wealth by reducing s ; a move that must be anticipated by the insurer's choice of π , triggering a race to the bottom.

¹⁵ Ruin theory predicts that if insurers only collected actuarially fair premiums, they eventually go bankrupt with certainty. Hence, a strictly positive loading λ is necessary to help ensure that the insurer exists at the time policyholders need it. Moreover, the qualifier "fair" is a technical convention established in the economics literature. It does not imply a value statement. Other definition coexist. For example, principle 4 of the Casualty Actuarial Society's Statement of Principles Regarding Property and Casualty Insurance Ratemaking specifies: "A rate is reasonable and not excessive, inadequate, or unfairly discriminatory if it is an actuarially sound estimate of the expected value of *all* future costs associated with an individual risk transfer." (emphasis added)

aversion σ and the shape of the defense function D . In general, higher loadings λ require more risk averse firms (σ) and less effective risk mitigation (β) for an insurance market to exist.

4.3 Information Asymmetries

In practice, firms are not homogeneous. Each firm i has an individual level of initial wealth w_i , risk aversion σ_i , and a defense function D_i .¹⁶ Consequently, each firm finds a different optimal action $(s_i, x_i)_{\sigma_i}^*$ from solving Equation (6). Information asymmetries arise because individual actions are not easily observable.

4.3.1 Adverse Selection

Adverse selection occurs before an insurance the contract is signed. The insurer faces the problem of finding the right premium π . For all choices of λ reasonable in the homogenous case, any choice of π in

$$\min_i (D_i(s_i^*) \cdot (1 + \lambda)) < \pi < \max_i (D_i(s_i^*) \cdot (1 + \lambda)) \quad (8)$$

leads to an adverse selection problem (Rothschild and Stiglitz, 1976). Firms can be divided in *good risks*, where π is high enough for the insurer to pay losses and profit, and *bad risks*, where π is too low for the individual probability of loss. Since the decision to seek insurance is made by rational firms, bad risks will seek (disproportionally more) insurance whereas good risks substitute insurance with security investment or risk acceptance. The resulting portfolio becomes a loss source for the insurer.¹⁷

¹⁶ In principle, also the loss size l may differ between firms. We stick with our approach to modify the loss distribution on the risk arrival side only in order to keep the equations simple. A practical model for numerical analysis should allow for more realistic loss distributions than Bernoulli trials.

¹⁷ Risk aversion and imperfect information on the demand side may counterbalance part of this problem in practice. Firms who do not know their risk exposure or security level, or who are highly risk averse, may stay in the pool despite being good risks in objective terms.

The canonical response against adverse selection is premium differentiation. In the best case, individual premiums π_i can be adjusted to the individual risk $D_i(s_i^*)$, where s_i^* may be anticipated. To simplify analytical models, D is often fixed for all firms, an assumption justifiable by the use of commodity technology (Carr, 2003). This allows the insurer to offer a menu of bundles (s_k, π_k) , such that each firm i chooses and commits to a security investment s_k in exchange for receiving coverage of the residual risk at premium π_k . In practice, security investment and actual security level are not perfectly aligned. Therefore, insurers have to resort to approximations of s_k by evaluating indicators on the firm’s technical and non-technical risk management practices.

4.3.2 Moral Hazard

Moral hazard occurs if the insurer cannot monitor policyholders’ contractual behavior. Once a firm has bought coverage $x = l$ at premium $\pi_j \cdot x$, it can profit from reducing its security investment to a value $s < s_k$ below the agreed level. If the insurer cannot observe the contract violation, it faces losses. Schwartz and Sastry (2014) note that insurers can only tolerate a limited share of such “malicious” policyholders in their portfolio.

Moral hazard can be dealt with by vigilant contract monitoring, possibly involving intermediaries such as technical auditors. Since cyber risk analysis is a complex and barely standardized task, the transaction costs caused by contract monitoring are rather high and might render certain types of insurance uneconomical. Another way to limit moral hazard are *deductibles* which technically define an upper bound for $x \leq x_{\max} < l$. The deductible $l - x_{\max}$ helps to better align the incentives of policyholders with the interest of the insurer (and by extension all policyholders in its pool).

4.3.3 Insurance Fraud

Insurance fraud occurs at the time of indemnification. Analyzing causal relationships, attributing losses to events, and verifying insurance policy conditions in (re-)programmable networked systems is a complex and time-consuming task. Digital forensics are in principle mature enough to provide reliable results after outside attacks. Nevertheless, ambiguities remain for technical and political reasons; the latter, if the attack involves a state actor, and diplomatic relations are concerned. Digital forensics reaches its limits, if an inside attacker falsifies traces or plants misleading evidence (Böhme et al., 2009). This may open a window of opportunity for fraudsters to overstate losses until the level of insurers' fraud prevention and detection systems in the claims management processes are comparable to those in conventional insurance. The opposite situation, where policyholders forgo a fraction of the loss which they cannot substantiate, acts as a barrier to the development of insurance markets (Ögüt et al., 2011).

4.4 Interdependent Security

While adverse selection and moral hazard are known and understood in conventional insurance, interdependent security is highly specific to the nature of cyber risk.¹⁸ Interdependent security is a special kind of externality best modeled by modifying the defense function: for each node in a network, the probability of a loss does not only depend on its own security (controls), but also on the security of all connected nodes. To illustrate this point, recall that the much-quoted Target data breach in late 2013 was caused by a security hole at one of its suppliers.

¹⁸ The term has been coined by Kunreuther and Heal (2003), while Varian (2002) has modeled the same phenomenon for cyber security slightly earlier.

4.4.1 Model and Example

Taking this to the level of firms in our economic model, the modified defense function I with interdependent security takes as arguments the investment of all connected firms:

$$p_i = I((s_1, \dots, s_n), g, i) = 1 - \prod_{j=1}^n (1 - g_{i,j} \cdot D(s_j)) . \quad (9)$$

In this specification, n is the number of firms and g is a $n \times n$ adjacency matrix of a (possibly directed) graph describing the first-order dependence between the nodes controlled by the firms. The matrix can hold binary values for all-or-nothing dependence or values between zero and one to weigh the importance of links. We require that $g_{ii} = 1 \forall i$ to generalize the case without interdependent security. Function I replaces D in Equations (5), (6), and (7).

To study the implication of interdependent security, it is convenient to regard symmetric interdependence between two firms 1 and 2 only. We obtain

$$p_1 = I(s_1, s_2, \alpha) = 1 - ((1 - D(s_1))(1 - \alpha \cdot D(s_2))) , \quad \text{and} \quad (10)$$

$$p_2 = I(s_2, s_1, \alpha) = 1 - ((1 - D(s_2))(1 - \alpha \cdot D(s_1))) , \quad (11)$$

where $\alpha \in]0, 1]$ is the degree of interdependence corresponding to the values in $g_{1,2}$ and $g_{2,1}$ of the general model. Without loss of generality we restrict the security investment to $s_i \in \{1, 2\}$, assume the security productivity $\beta = 1/2$, initial wealth $w = 3$, and the loss $l = 4$. Table 1 shows the payoffs of a matrix game between both firms deciding on security investment without the option to seek market insurance as a function of α .¹⁹

For any degree of interdependence $\alpha > 0$, the only pure strategy Nash equilibrium is $(s_1, s_2) = (1, 1)$ – both firms invest little in security – because no firm can improve its

¹⁹ The example assumes risk neutral firms because we do not consider market insurance. Qualitatively similar results can be found for models with risk aversion (e. g., Johnson et al., 2011).

Table 1: Interdependent cyber risk as a matrix game between two defenders

Security of firm 1	Security of firm 2	
	low: $s_2 = 1$	high: $s_2 = 2$
low: $s_1 = 1$	$-\alpha, -\alpha$	$-\frac{1}{2}\alpha, -\frac{3}{2}\alpha$
high: $s_1 = 2$	$-\frac{3}{2}\alpha, -\frac{1}{2}\alpha$	$-\frac{3}{4}\alpha, -\frac{3}{4}\alpha$

payoff by unilaterally switching to $s_i = 2$. However, the social optimum of this game is $(s_1, s_2) = (2, 2)$ – both firms invest much in security. This indicates a classical prisoner’s dilemma. In other words, if networked firms secure their nodes, interdependent nodes may free-ride on their efforts. As a result, no single party wants to contribute to the network security, which exhibits characteristics of a public good in this kind of model.

4.4.2 Implications

Interdependent security has two important implications for insurers. First, insurers not only face the difficulty of measuring and monitoring the security practices and exposure of their policy holders. To precisely model the risk arrival, they also need to collect information about the security of interconnected nodes. These may be owned and operated by parties without contractual relationship with the insurer, possibly residing in different jurisdictions. In the absence of this information, insurers must resort to conservative approximations. They may at least consider to include the number of incoming edges and indicators about the degree of interdependence in the premium calculation.

Second, insurers may provide the social coordination mechanism to resolve this instance of the prisoner’s dilemma. If the insurer can observe firms’ security effort, it can (and should) design a pricing scheme which incentivizes security investment at the social optimum. This is in the best interest of all policyholders and leads to welfare improvements

beyond of what is achievable with risk pooling in conventional insurance lines. This idea is a strong argument for bootstrapping a deep market for cyber risk insurance, possibly with the help of subsidies justifiable with the provision of a public good. However, important limitations remain. There may be free-riders who do not buy insurance. They enjoy the benefits of the public good while having no incentive to contribute. Another issue is competition in the insurance market. A monopolist is in the best position to act as social coordinator, but it is well-known that absence of competition leads to inefficiencies and premiums levels that generate monopoly rents. However, if multiple carriers compete and insure different parts of interdependent networks, they would have to engage in information sharing at a depth and speed unprecedented in conventional lines of insurance. To thwart the concerns of price-rigging in such close relationships between competitors, the whole scheme would have to be run under the watchful eyes of a trusted party.

4.4.3 Literature

Interdependent security has inspired many analytical models.²⁰ Most prominently, Ögüt et al. (2005) are the first to combine the concept with market insurance and study the social coordination mechanism. Hofmann (2007) considers imperfect information and adds premium differentiation to the analysis. Bolot and Lelarge (2008) explore mean-field approximation to analyze the possibility of coordinated solution with the help of insurance for selected network topologies (also Lelarge and Bolot, 2009). Shetty et al. (2010) note that a coordination mechanism that maximizes welfare with the help of market insurance will not, in general, maximize the overall security investment (later extended and refined

²⁰ Many economic analyses of cyber risk insurance have appeared in technical venues. This has led to some abuse of terminology, which may cause confusion to people entering the area. For example, a conference paper with “market analysis” in its title does not model a market (Pal et al., 2014); or a workshop paper with “insurance” in its title does not consider risk aversion (a necessary condition for insurance), but models a social redistribution scheme (Naghizadeh and Liu, 2014); Ögüt et al. (2011) use “correlation” in the title but model interdependent risk; Shetty et al. (2010) model an unconventional type of risk aversion by leaving the security investment outside the utility function (unlike Equation (27) of Ehrlich and Becker (1972)).

in Schwartz and Sastry, 2014). Johnson et al. (2011) study symmetric equilibria in fully connected graphs where firms can buy market insurance and invest in two types of security controls, one that generates externalities and another one that does not. Against the backdrop of this research, it is remarkable that we could not find any empirical quantification of interdependent security in a real network.

4.5 Market Structure

The insurance business benefits from economies of scale and scope. The amount of safety capital needed per policyholder is smaller for large and diverse portfolios for several reasons. The central limit theorem predicts that realized losses deviate less from the distribution mean as the portfolio size increases, and fixed administrative costs can be distributed over many policyholders. Moreover, transaction costs per contract (for underwriting and monitoring) are smaller if each policyholder has multiple policies from the same insurer. Empirical evidence suggest that some conventional lines of insurance have features of a natural monopoly (Emons, 2001), although the data confounds public and privately run carriers (e. g., von Ungern-Sternberg, 1996). Cyber risk insurance generates additional scale effects related to the monitoring of policyholders and the topology of their connections to interdependent systems owned and operated by others. Finally, the social coordination function (see Section 4.4.2) would require tight guidance and widely accepted standards if it is to be realized by competing organizations.

Another aspect is the market structure of potential policyholders. Insurance is most efficient if insurers can pool many diversifiable risks, each big enough to be substantial for the policyholder, but not catastrophic for the insurer. Concentration among policyholders renders insurance less attractive because firms may prefer to accept small risks, avoiding the transaction costs of risk transfer. Risk acceptance remains problematic for larger or internally correlated risks (Böhme and Kataria, 2006), but there may remain too few

policyholders to pool such risks. This is one facet of the more general “ $N = 1$ problem” in the technology sector, where dominant suppliers are unparalleled. Tendencies of such concentration can be observed in industries such as standard software, Internet search, cloud hosting, and online social networking.

Finally, how does the existence of a market for cyber risk insurance affect competition? Arguably, if insurance can price cyber risk precisely, then premium differentiation would make more diverse risks less expensive to insure. This could partially offset the advantage of incumbent platforms and therefore stimulate technical diversity and economic competition. Both are considered desirable goals that generate positive externalities on firms who do not participate in a market for cyber risk insurance.

5 Actuarial Modeling

Formal economic models often simplify loss distributions to central moments or Bernoulli trials. By contrast, premium calculation in practice requires knowledge of the full loss distribution in order to determine the right amount of safety capital (cf. Equation (7)).

5.1 Data Sources

The preferred data source for actuarial modeling in conventional insurance are historical claims. Due to the novelty of cyber risk, the nascent insurance market cannot look back at a long history. Deductibles and rational under-reporting keep the number of data points artificially low. Moreover, historical data is of limited use, because threats, vulnerabilities, and controls evolve rapidly with the pace of the technological development.

Technical indicators are an alternative data source for monitoring the threat landscape. Typical measurement setups include

- sensors measuring traces of known attacks (e. g., packet backscatter or passive DNS

analysis, see Bilge et al. (2011)),

- spam traps and “honeypot” computers that mimic vulnerable devices and are instrumented to sense attack behavior (Provos and Holz, 2007), as well as
- trackers observing command-and-control information of the botnets²¹ that serve as primary infrastructure for cybercrime.

Several organizations collect this type of data and make it available to interested parties.²² Proprietary threat intelligence may also include information collected by vendors of antiviral software on their clients’ devices. A shortcoming of these data sources is that they do not include controls and ultimately loss amounts.

The security practices of firms (and to some extent their suppliers) are primarily measured with self-assessments, typically as part of the underwriting process. Moreover, a new industry of intermediaries for cyber risk analysis is forming. These services promise security ratings based on remote technical measurements (e. g., scans of the corporate Internet address space, maturity of the corporate website technology, etc.) and aggregation of incident reports.²³ In contrast to self-assessments, rating availability does not depend on the cooperation of the rated firms. However, due to proprietary models and the novelty of these businesses, little is known about the accuracy of the ratings and their robustness to attempts of deception.

With the introduction of breach reporting obligations for the loss of personal data, public and private registers began to collect information on incidents at the firm level.²⁴ Incidents

²¹ A *botnet* (robot network) is a network of compromised computers, reprogrammed to be controlled remotely by criminals.

²² For example <https://zeustracker.abuse.ch> by the feed providers Zeus Tracker, <http://www.confickerworkinggroup.org> by The Conficker Working Group, <https://www.shadowserver.org> by the Shadowserver Foundation, <https://honeynet.org> by The Honeynet Project, or <http://www.zone-h.org> that was created by Roberto Preatoni.

²³ For example the enterprises SecurityScorecard, Bitsight, and QuadMetrics.

²⁴ For example, <https://datalossdb.org> operated by the Open Security Foundation, <http://www.privacyrights.org/data-breach> by the Privacy Rights Clearinghouse – see Edwards et al. (2016)

at the device (or network address) level can be inferred from various blacklists which are compiled from end users' abuse reports as well as from technical measurements. Law enforcement agencies maintain registers of cyber-related incidents primarily to compile official police statistics. However, these sources are known to be particularly prone to under-reporting. We expect that registers will become more relevant in the future as many jurisdictions plan mandatory reporting for general security incidents and sanction non-compliance (e. g., EC, 2013).

Victimization surveys²⁵ complement the list of data sources. Unlike technical indicators and incident reports, surveys draw a more representative picture of the impacts of cyber attacks on consumers and businesses. They may serve as basis to estimate the baseline risk. Well-known limitations of consumer surveys are that questions must be notoriously superficial and the resulting data is prone to response biases (Florêncio and Herley, 2011).

The last data source are indicators inspired from finance. For example, Geer and Pareek collect and publish a monthly sentiment indicator from a panel of security experts.²⁶ Others suggest to monitor underground marketplaces for prices of goods and services related to cyber crime (Thomas et al., 2015; HPE, 2016) or track volumes and prices on markets for vulnerabilities (Zhao et al., 2015).

5.2 Individual Loss Distribution

Research on loss distributions for cyber risk is scattered over multiple disciplines. The shortage of data may require Bayesian methods for extracting the maximum information from a few data points and finding appropriate ways of incorporating expert knowledge.

for a statistical model of breach sizes in this dataset – , and the VERIS Community Database <http://vcdb.org>, one of the data sources of the annual data breach investigation report published by Verizon.

²⁵ Prominent examples are the annual reports of the Computer Security Institute (CSI) that got published between 1996–2011 (e. g., CSI, 2011), periodical Eurobarometer studies (e. g., EC, 2015), and global studies carried out by private firms, (e. g., PwC, 2016).

²⁶ <http://www.cybersecurityindex.org>

Actuaries also need to find the right balance between parsimonious models and representative models, which convey the complexity of the networked environment. However, there is no consensus on data sources or appropriate methodology so far. Industry reports often provide estimates of central moments. NetDiligence (2016), for example, analyze 176 cyber insurance claims. They report an average cost per incident of about \$0.7 million in 2016. In a study on insurability of cyber risk, Biener et al. (2015) select 994 cyber events from a commercial database of publicly reported general operational losses between 1971 and 2009. They report a mean loss per cyber incident of \$40.5 million. Several authors (Gordon et al., 2011; Gatzlaff and McCullough, 2010) consistently report a negative stock market reaction following the announcement of breaches at publicly traded companies. The effect is in the order of 1% in terms of short-term cumulative abnormal return, a common metric of the event study method (MacKinlay, 1997). This translates to losses per incident in the order of \$300 million under the assumption of an average market capitalization of NASDAQ-100 companies between 2000 and 2010.²⁷

All these figures are not directly comparable because contexts, conditions and methods vary between studies. Should average losses be reported per record, per incident, per firm, or for entire sectors or economies? Shall losses include indirect costs? Shall they include or exclude recovered values? Unfortunately, some of the frequently cited studies are not even transparent about their methodology. Academics have flagged these issues (e.g., Anderson et al., 2013; Riek et al., 2016), but a commonly agreed knowledge base is still not within reach. One may hope that the experience of the insurance industry will facilitate the development of comparable standards for this risk class.

Even if estimates of central moments were accurate and comparable, they reveal very little about the shape of a loss distribution. Most loss distributions are skewed to the right.

²⁷ The number is considerably smaller at \$80 million for the average company in the NYSE Composite, which includes smaller and less technology-focussed companies.

For example, the mean estimate of \$ 40.5 million in Biener et al. (2015) originates from a distribution with a median of only \$ 1.9 million. Its mean is inflated by a few outliers of up to \$ 13.3 billion. This calls for a closer look at the tails of the individual loss distributions. In comparison to operational losses not classified as cyber, Biener et al. (2015) observe that cyber losses exhibit shorter (right) tails than conventional risks. This somewhat qualifies concerns raised by earlier studies reporting distinctive features of heavy tails in cyber loss data (Maillart and Sornette, 2010; Edwards et al., 2016; Wheatley et al., 2016).²⁸

The left tail deserves attention as well. Data from a consumer victimization survey indicates that loss distributions are zero-inflated even after controlling for the (rather unlikely) event of becoming a victim of cyber crime. According to Riek et al. (2016), only 33% of identity theft²⁹ victims in online banking suffer monetary losses. This subset of victims experiences direct financial losses of on average €2,150 per incident; the median is €630. This highlights the errors caused by replacing these kind of loss distributions with a simple Bernoulli loss model (see Section 4.2), where the probability of a fixed loss of €2,150 would be $p = 0.007$ in five years, suggesting an actuarially fair (i. e., lower bound) premium of about €3 per year if the entire population was insured.³⁰ Recall that knowledge of the shape of the individual loss distribution is necessary (but not sufficient) for determining the compound loss distribution.

²⁸ Note that heavy tails are observed primarily for the count data of breached records whereas moderate tails are observed in monetary losses (of otherwise incomparable data sources). Besides many other factors, this may hint to a concave relationship between the number of breached records and the total cost of a breach.

²⁹ Attackers impersonate their victims using stolen access credentials or identifying information.

³⁰ The survey asked for incidents in the past five years. About 2% of the sample reported victimization experience for the specific case of identity theft in online banking. The data is representative for the adult population of Internet users in six European countries: Germany, Estonia, Italy, The Netherlands, Poland, and the UK.

5.3 Compound Loss Distribution

The compound loss distribution determines the realization of aggregate claims against an insurer who covers a pool of many risks. It is given by the convolution of all individual loss distributions if the risks are statistically independent. However, this important condition does not apply to large parts of cyber risk. The very success factors of information technology, economies of scale, standardization of programmable devices, and global networks, create complicated dependence structures between loss distributions of networked components. This is because standard software has common vulnerabilities that can, in the worst case, be exploited remotely at global scale.

The risk of correlated failure has been prominently mentioned by Geer et al. (2003), at the time attributed to the near-monopoly of Microsoft on the desktop software market. With Unix-based alternatives regaining strength as well as competing and fragmented platforms in the mobile market, the situation has somewhat changed. However, vulnerabilities in popular libraries such as zlib or OpenSSL, which are open source and widely used on all platforms, remain a source of concern. Similarly, concentration on the market for basic cloud services (infrastructure and platforms) may lead to correlated failures of downstream cloud services. Consequently, an insurer's profitability in cloud services policies (or business interruption if the business depends on such services) lasts on the second or third-tier cloud operators' ability to avoid single points of failure.³¹

Positively correlated risks generally shift the probability mass of the compound distribution to the extremes, thereby thwarting the balancing property of pooled risk, the very principle that makes insurance economically viable. Figure 5 illustrates this for $n = 200$ homogeneous Bernoulli risks, each causing a fixed loss of $l = \$1$ million with probability $p = 0.2$. We use a single-factor model of correlation, where a fraction ρ of the loss

³¹ Insurers may exclude this type of risk, but the resulting policies will be considerably less attractive to businesses using cloud services. They might also try to hold suppliers liable after incidents, but large events will soon exhaust the suppliers' capital stock.

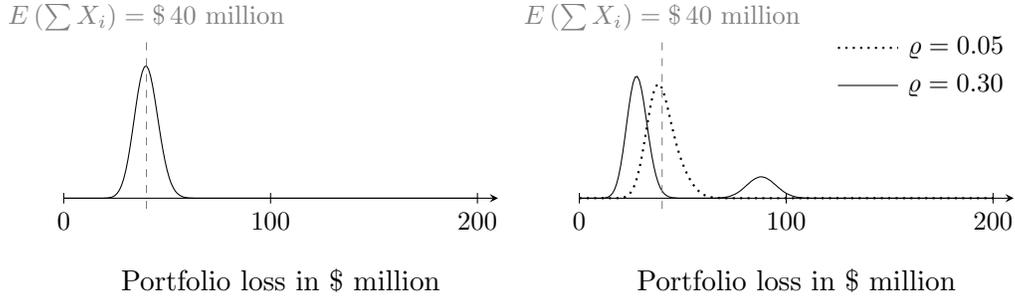


Figure 5: Probability density of compound loss distribution for uncorrelated (left) and correlated (right) risks

probability is determined by a common Bernoulli trial, for instance a newly discovered vulnerability in standard software. The model is set up to keep the expected portfolio loss independent of the strengths of correlation ρ (see Böhme, 2005, for the specification). It demonstrates that correlation affects the shape of the compound loss distribution. With increasing ρ , more and more probability mass is shifted to the sensitive right tail. In the extreme case of $\rho = 0.3$, the insurer faces a non-negligible risk of losing \$100 million in a period, more than twice as much as he can expect in premiums if the actuarial fair premium is taken as a baseline. To some extent, this can be compensated by increasing the loading λ (cf. Equation (7)), but the market disappears if higher premiums π render risk transfer uneconomical for firms' levels of risk aversion.

Estimating the compound loss distribution for a portfolio of real cyber risks is a major challenge. Cyber risk insurance suffers from a lack of reliable data in general, and estimating multivariate properties, such as correlation coefficients between distributions, demands even more data points than estimates of univariate (central) moments. Moreover, the dependence structure between many risks is most likely not homogenous, and may be more complex than linear correlation between central moments. Extreme value theory has been applied in finance and insurance (Embrechts et al., 1999) in order to characterize and estimate tail risk in compound loss distributions. If individual risks are dependent, copulas

offer very general models of dependence between variables that can be parametrized to account for co-occurrence of extreme events (Nelsen, 1999). These ideas have been brought to cyber risk analysis. Using a t -copula, which captures tail dependence, Böhme and Kataria (2006) report evidence for positive dependence between the attack activity observed at 35 globally distributed sensors (honeypots, cf. Section 5.1). This data source tracks threat information, not losses or claims. The result can nevertheless be interpreted as an indication of structural dependence, which will also drive the relevant loss distributions for insurers.³²

In the absence of empirical data, analysts must resort to simulations. They can leverage the network environment component of the cyber risk insurance framework (Figure 3) to study propagation of risk along the edges of the graph structure. Lorenz et al. (2009) propose a methodology based on the popular (but often misleading) mean field approximation.³³ Johnson et al. (2014) study the computational complexity of the exact solution for different classes of topologies. While homogeneous networks and star-shaped topologies (see the second and third example in Figure 4) permit efficient solutions, the problem is NP-hard in general. For the special case of scale-free networks, the authors show that the compound loss distribution cannot be computed from a random sample of nodes, a plausible situation for an insurer who shares the market with competitors or firms who mitigate, accept or avoid their risks. Many natural networks, foremost the routing topology of the Internet, exhibit scale-free characteristics.

Note that this topology may be, but is not necessarily aligned with the topology governing interdependent security (see Section 4.4). Both concepts are related, but distinct. Interdependent security connects defense functions and thus parameters of individual loss distributions. By contrast, risk propagation leads to cumulated risk, that is dependence

³² Copulas have also been proposed to obtain more precise *individual* loss distributions of large enterprises where a single threat can affect many devices at the same time (see Herath and Herath, 2007).

³³ The drastic simplification of this approach drops all local information about interactions between nodes and replaces it with a global average.

between realizations of losses. Correlation is a special form of the latter. In other words, firms anticipate interdependent security in their cyber risk management, whereas insurers (and their regulators) must primarily be concerned about cumulated risk.

The fear of cumulated cyber risk led reinsurers to explicitly exclude cyber risk in the early 2000s, a fact still named as barrier to the development of cyber risk insurance a decade later (ENISA et al., 2012). Alternative forms of bulk cyber risk transfer, such as securitization or catastrophe bonds have been discussed, but cause concerns about moral hazard by investors who may find themselves in a position where they benefit from cyber incidents (Anderson et al., 2008).

6 Conclusion

Cyber risk insurance as a risk management tool has not kept pace with the adoption of information technology. At the time of writing, the market is about to develop, primarily in comfortable niches, but increasingly taking on more substantial exposure. If there is anything to be learnt from the last financial meltdown, adding layers of indirection is only a good idea if the party who takes the risk is in a better position to understand, mitigate, and eventually bear it. Cyber risk quantification faces many challenges in the first place. It differs, as laid out in this article, in many important respects from the conventional risks that the insurance industry has built an expertise to model and price.

This article could close with the known mantra and call for more actuarial data; or complain about the chicken-and-egg problem that actuarial data will not become available unless policies are written and claims filed. However, even decades of claims about “yesterday’s attacks” will not inform about “tomorrow’s risk” in a domain where “moving-target defense” has become a dictum. Models of cyber risk arrival need to be more predictive. They must draw on the available data at earlier stages of our cascade model of cyber risk

arrival (cf. Figure 2). This includes information about the dependence topology. This calls for data sharing frameworks between competitors or with intermediaries in order to be able to calculate compound loss distributions for the parts of a larger network in an insurer’s portfolio.³⁴ Some authors see a role for the government as a standard-setter to facilitate this exchange (e. g., Biener et al., 2015).

More systematic collection and evaluation of data at early stages is also essential for identifying risk factors and quantifying their effects. The checklists used for underwriting cyber risk today may help lawyers identifying conditions precedent to liability, but they seem too superficial for meaningful premium differentiation. Premium differentiation, however, is a crucial instrument to stimulate better security practices at deployment and operations and, in the long run, pass the signal upward the supply chain, where it could stimulate more risk-conscious software development and systems engineering (Heitzenrater et al., 2016). If this channel is too indirect and slow, insurers could take a more active role and liaise with software vendors. Laszka and Grossklags (2015) outline how insurers can proactively try to remove software vulnerabilities in order to reduce correlated risk and tighten the critical right tail of the compound loss distribution. While the idea is appealing, and has been mentioned by Anderson (2008) and Böhme (2005), it implies that insurers enter an unfamiliar partnership with the software industry. Moreover, it requires collective action because no insurer has incentives to be the first mover if all competitors benefit equally from the fruits of the efforts.³⁵ These and other lessons learned from the analysis in this article are summarized in Box 3.

³⁴ Compliance with privacy and data protection laws is an important requirement for such information sharing arrangements. The legal and technical aspects thereof require further scrutiny.

³⁵ A conceivable, but morally and economically questionable alternative would be to allow for hardened versions of software or services for policyholders of “activist insurers”. The increasing system diversity would nonetheless generate positive externalities for competitors and uninsured firms.

Box 3: Key messages

- We propose a differentiated view on cyber versus conventional risk by separating the nature of risk arrival from the target exposed to risk.
- Cyber risk is *technically* characterized by high design complexity, (re)programmable behavior of networked components, and a global dynamic threat surface.
- Cyber risk is *economically* characterized by incomplete information, externalities, and correlation caused by common risk factors.
- Cyber security is a timing game of information on threats and vulnerabilities.
- Cyber risk insurers need to establish communication channels with policyholders in order to quickly react to changes in the threat environment.
- Quantification of cyber risk suffers from a lack of relevant data. This is due to missing standards for data collection and missing incentives for data sharing.
- Cyber risk analysis should emphasize early indicators over historical claims. It can draw on methods of network science to estimate portfolio loss distributions.
- Cyber risk insurers may need to share more information with competitors than in conventional lines of insurance.
- Cyber risk and insurance are nascent fields of interdisciplinary research.
- As information technology gets pervasive, more industries follow the model of the ICT industry. More risks will take the characteristics of cyber risk.

Foreseeable future developments include that more industries will follow the model and the economic logic of the software industry. Consequently, more risk is taking the characteristics of cyber risk. Current technology trends, such as the Internet of Things, the wake of a fourth industrial revolution, autonomous vehicles, and mass-customized medication based on computer-assisted diagnosis, just to name a few, support this point. In this context, the need for cyber risk analysis is paramount on the individual and compound level. It is barely responsible to embrace further increases in the society's dependence on information technology without being able to monitor the resulting cyber risk with principled

and scientific methods.

Another stream of current technology development explores massively distributed systems formed by pseudonymous actors who execute long-running protocols on top of the Internet infrastructure. These protocols update a cryptographically secured global data structure and resolve inconsistencies and conflicts without a central party. The success of Bitcoin, a virtual currency scheme and the most prominent example today, has drawn attention to this technological paradigm (Böhme et al., 2015). Although its future is subject to high uncertainty, a possible wider adoption has two important implications for cyber risk analysis. First, distributed ledgers (an umbrella term for said systems) produce a very different sort of cyber risk than described in this article. Besides technological factors, this “novel” cyber risk is determined by new “laws of nature” enforced by cryptography as well as the aggregate behavior of many autonomous algorithms reacting to these laws. Ground-breaking research and modeling effort is needed before these risks can be priced and transferred. Second, this technology has become a platform for disruptive innovation of financial intermediation, with Bitcoin taking the lead on payments systems. Speculatively, this may change not only the type of risk insured, but revolutionize the way we think about and organize insurance.

Acknowledgements

The authors thank Avraham Adler and Karen Sonnet for the comments and suggestions provided on earlier versions of this report and the anonymous reviewers for their constructive feedback, which significantly improved the final version.

References

- Alessandro Acquisti, Allan Friedman, and Rahul Telang. Is there a cost to privacy breaches? An event study. In *Workshop on Economics of Information Security (WEIS)*, University of Cambridge, UK, 2006.
- AGCS. A guide to cyber risk: Managing the impact of increasing interconnectivity. Technical report, Allianz Global Corporate & Specialty (AGCS), 2015.
- Ali Alkaabi, George Mohay, Adrian McCullagh, and Nicholas Chantler. Dealing with the problem of cybercrime. In Ibrahim Baggili, editor, *Digital forensics and cyber crime*, volume 53 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 1–18, Berlin, Heidelberg, 2011. Springer.
- Ross Anderson. *Security engineering: A guide to building dependable distributed systems*. Wiley, 2 edition, 2008.
- Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314 (5799):610–613, 2006.
- Ross Anderson, Rainer Böhme, Richard Clayton, and Tyler Moore. Security economics and the internal market. Technical report, European Network and Information Security Agency, ENISA, 2008.
- Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J. G. van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In Rainer Böhme, editor, *The economics of information security and privacy*, chapter 12, pages 265–300. Springer, Berlin, Heidelberg, 2013.
- Ashish Arora, Rahul Telang, and Hao Xu. Optimal policy for software vulnerability disclosure. *Management Science*, 54(4):642–656, 2008.

- Hadi Asghari, Michael Ciere, and Michel J. G. van Eeten. Post-mortem of a zombie: Conficker cleanup after six years. In *USENIX Security Symposium*, Washington, D. C., USA, 2015.
- Walter Baer. Rewarding IT security in the marketplace. *Contemporary Security Policy*, 24(1):190–208, 2003.
- Magdalena Balazinska, Bill Howe, and Dan Suciu. Data markets in the cloud: An opportunity for the database community. In H. V. Jagadish, editor, *Proceedings of the Very Large Database Endowment (PCLDB)*, volume 4, pages 1482–1485, Seattle, WA, USA, 2011.
- Tridib Bandyopadhyay, Vijay S. Mookerjee, and Ram C. Rao. Why IT managers don’t go for cyber-insurance products. *Communications of the ACM*, 52(11):68–73, 2009.
- Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, and Chris Wright. Timing the application of security patches for optimal uptime. In *Proceedings of the USENIX Systems Administration Conference (LISA)*, pages 233–242, Philadelphia, PA, USA, 2002.
- Gary S. Becker. Crime and punishment: An economic approach. In Gary S. Becker and William M. Landes, editors, *Essays in the economics of crime and punishment*, pages 1–54. National Bureau of Economic Research, 1974.
- Richard S. Betterley. Cyber/privacy insurance market survey – 2016. Technical report, Betterley, 2016.
- Christian Biener, Martin Eling, and Jan Hendrik Wirfs. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance—Issues and Practice*, 40(1):131–158, 2015.

- Leyla Bilge, Engin Kirda, Christopher Kruegel, and Marco Balduzzi. EXPOSURE: Finding malicious domains using passive DNS analysis. In *Proceedings of the Annual Network & Distributed System Security Symposium (NDSS)*, San Diego, CA, USA, 2011.
- Rainer Böhme. Cyber-insurance revisited. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, USA, 2005.
- Rainer Böhme. A comparison of market approaches to software vulnerability disclosure. In Günter Müller, editor, *Emerging Trends in Information and Communication Security (ETRICS)*, volume 3995 of *Lecture Notes in Computer Science*, pages 298–311. Springer, 2006.
- Rainer Böhme. Security metrics and security investment models. In Isao Echizen, Noboru Kunihiro, and Ryoichi Sasaki, editors, *Advances in Information and Computer Security*, volume 6434 of *Lecture Notes in Computer Science*, pages 10–24, Berlin, Heidelberg, 2010. Springer.
- Rainer Böhme and Gaurav Kataria. Models and measures for correlation in cyber-insurance. In *Workshop on the Economics of Information Security (WEIS)*, University of Cambridge, UK, 2006.
- Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, USA, 2010.
- Rainer Böhme, Felix C. Freiling, Thomas Gloe, and Matthias Kirchner. Multimedia forensics is not computer forensics. In Zeno J. M. H. Geradts, Katrin Franke, and Cor J. Veenman, editors, *International Workshop on Computational Forensics (IWCF)*, volume 5718 of *Lecture Notes in Computer Science*, pages 90–103, Berlin, Heidelberg, 2009. Springer.

- Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–238, 2015.
- Jean C. Bolot and Marc Lelarge. A new perspective on Internet security using insurance. In *Conference on Computer Communications (IEEE INFOCOM)*, Phoenix, AZ, USA, 2008.
- Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. The quest to replace passwords: A framework for comparative evaluation of Web authentication schemes. Technical Report 817, University of Cambridge Computer Laboratory, UK, 2012.
- Erik Brynjolfsson and Lorin M. Hitt. Computing productivity: Firm-level evidence. *Review of Economics and Statistics*, 85(4):793–808, 2003.
- Erik Brynjolfsson, Lorin M. Hitt, and Shinkyu Yang. Intangible assets: Computers and organizational capital. *Brookings Papers on Economic Activity*, pages 137–199, 2002.
- Richard Caralli, James Stevens, Lisa Young, and William Wilson. Introducing OCTAVE allegro: Improving the information security risk assessment process. Technical Report CMU/SEI-2007-TR-012, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.
- Nicholas G. Carr. IT doesn’t matter. *Harvard Business Review*, May:5–12, 2003.
- Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. A model for evaluating IT security investments. *Communications of the ACM*, 47(7):87–92, 2004.
- James J. Cebula, Mary E. Popeck, and Lisa R. Young. A taxonomy of operational cyber security risks. Technical report, Software Engineering Institute, 2010.

- Thomas M. Chen. Stuxnet, the real start of cyber warfare? *IEEE Network*, 24(6):2–3, 2010.
- Richard A. Clarke and Robert K. Knake. *Cyber war: The next threat to national security and what to do about it*. Harper Collins, 2012.
- Donald R. Cressey. *Other people’s money; a study of the social psychology of embezzlement*. Free Press, Glencoe, IL, USA, 1953.
- CSI. 2010/2011 computer crime and security survey. Technical report, Computer Security Institute (CSI), 2011.
- EC. Proposal for a Directive of the European Parliament and of the Council concerning measures to ensure a high common level of network and information security across the Union. *COM (2013) 48 final*, 2013.
- EC. Special Eurobarometer 423: Cyber security. Technical report, European Commission (EC), 2015.
- Benjamin Edelman. Adverse selection in online “trust” certifications and search results. *Electronic Commerce Research and Applications*, 10(1):17–25, 2011.
- Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity*, 2(1):3–14, 2016.
- Isaac Ehrlich and Gary S. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, 1972.
- Martin Eling and Werner Schnell. Ten key questions on cyber risk and cyber risk insurance. Technical report, The Geneva Association, 2016.
- Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling extremal events: For insurance and finance*. Springer, Berlin, Heidelberg, 2 edition, 1999.

- Winand Emons. Imperfect tests and natural insurance monopolies. *The Journal of Industrial Economics*, 49(3):247–268, 2001.
- ENISA, Neil Robinson, and RAND Europe. Incentives and barriers of the cyber insurance market in Europe. Technical report, European Union Agency for Network and Information Security (ENISA), 2012.
- Dinei Florêncio and Cormac Herley. Sex, lies and cyber-crime surveys. In *Workshop on Economics of Information Security (WEIS)*, George Mason University, Fairfax, VA, USA, 2011.
- Kevin M. Gatzlaff. Implications of privacy breaches for insurers. *Journal of Insurance Regulation*, 31:197–214, 2012.
- Kevin M. Gatzlaff and Kathleen A. McCullough. The effect of data breaches on shareholder wealth. *Risk Management and Insurance Review*, 13(1):61–83, 2010.
- Dan Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quarterman, and Bruce Schneier. CyberInsecurity: The cost of monopoly. How the dominance of Microsoft’s products poses a risk to security. Technical report, Computers & Communications Industry Association, 2003.
- William Gibson. *Burning chrome*. Ace Books, 1987.
- Marc D. Goodman. Why the police don’t care about computer crime. *Harvard Journal of Law & Technology*, 10(3):465–495, 1997.
- Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 2002.

- Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. The impact of information security breaches: Has there been a downward shift in costs? *Journal of Computer Security*, 19(1):33–56, 2011.
- Torsten Grzebiela. Insurability of electronic commerce risks. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS)*, Big Island, HI, USA, 2002.
- Chad Heitzenrater, Rainer Böhme, and Andrew Simpson. The days before zero day: Investment models for secure software engineering. In *Workshop on Economics of Information Security (WEIS)*, University of California, Berkeley, CA, USA, 2016.
- Hemantha S. B. Herath and Tejaswini C. Herath. Cyber-insurance: Copula pricing framework and implication for risk management. In *Workshop on Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.
- Cormac Herley. Security, cybercrime, and scale. *Communications of the ACM*, 57(9):64–71, 2014.
- Annette Hofmann. Internalizing externalities of loss prevention through insurance monopoly: An analysis of interdependent risks. *Geneva Risk and Insurance Review*, 32(1):91–111, 2007.
- Kevin J. S. Hoo. How much is enough? A risk management approach to computer security. In *Workshop on Economics of Information Security (WEIS)*, University of California, Berkeley, CA, USA, 2002.
- HPE. HPE security research: Monetizing stolen credit card data. Technical report, Hewlett Packard Enterprise (HPE), 2016.
- Christos Ioannidis, David Pym, and Julian Williams. Information security trade-offs and

- optimal patching policies. *European Journal of Operational Research*, 216(2):434–444, 2012.
- ISO/IEC. ISO/IEC 27000:2014: Information technology – Security techniques – Information security management systems – Overview and vocabulary. Standard, International Organization for Standardization/International Electrotechnical Commission (ISO/IEC), 2014.
- Benjamin Johnson, Rainer Böhme, and Jens Grossklags. Security games with market insurance. In John S. Baras, Jonathan Katz, and Eitan Altmann, editors, *Decision and game theory for security*, volume 7037 of *Lecture Notes in Computer Science*, pages 117–130, Berlin, Heidelberg, 2011. Springer.
- Benjamin Johnson, Aron Laszka, and Jens Grossklags. The complexity of estimating systematic risk in networks. In *Proceedings of the Computer Security Foundations Symposium (CSF)*, pages 325–336, Vienna, Austria, 2014.
- Stanley Kaplan and B. John Garrick. On the quantitative definition of risk. *Risk analysis*, 1(1):11–27, 1981.
- Keith Kirkpatrick. Cyber policies on the rise. *Communications of the ACM*, 58(10):21–23, 2015.
- Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2/3):231–249, 2003.
- Aron Laszka and Jens Grossklags. Should cyber-insurance providers invest in software security? In Günther Pernul, Peter Y. A. Ryan, and Edgar Weippl, editors, *Computer Security – ESORICS 2015*, volume 9326 of *Lecture Notes in Computer Science*, pages 483–502, Berlin, Heidelberg, 2015. Springer.

- Stefan Laube and Rainer Böhme. The economics of mandatory security breach reporting to authorities. *Journal of Cybersecurity*, 2(1):29–41, 2016.
- Marc Lelarge and Jean Bolot. Economic incentives to increase security in the Internet: The case for insurance. In *Conference on Computer Communications (IEEE INFOCOM)*, Rio de Janeiro, Brazil, 2009.
- Jan Lorenz, Stefano Battiston, and Frank Schweitzer. Systemic risk in a unifying framework for cascading processes on networks. *The European Physical Journal B*, 71(4):441–460, 2009.
- William J. Lynn. Defending a new domain. *Foreign Affairs*, 89(5):97–108, 2010.
- A. Craig MacKinlay. Event studies in economics and finance. *Journal of Economic Literature*, 35(1):13–39, 1997.
- Thomas Maillart and Didier Sornette. Heavy-tailed distribution of cyber-risks. *The European Physical Journal B*, 75(3):357–364, 2010.
- Gennady Medvinsky, Charlie Lai, and B. Clifford Neuman. Endorsements, licensing, and insurance for distributed system services. In *Proceedings of the ACM Conference on Computer & Communications Security (ACM CCS)*, pages 170–175, Fairfax, VA, USA, 1994.
- Robert Irwin Mehr and Emerson Cammack. *Principles of insurance*. Homewood, 5 edition, 1972.
- Charlie Miller. The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.

- Tyler Moore, Allan Friedman, and Ariel D. Procaccia. Would a ‘cyber warrior’ protect us: Exploring trade-offs between attack and defense of information systems. In *Proceedings of the Workshop on New Security Paradigms (NSPW)*, pages 85–94, Concord, MA, USA, 2010.
- Arunabha Mukhopadhyay, Samir Chatterjee, Debashis Saha, Ambuj Mahanti, and Samir K. Sadhukhan. Cyber-risk decision models: To insure IT or not? *Decision Support Systems*, 56:11–26, 2013.
- Parinaz Naghizadeh and Mingyan Liu. Voluntary participation in cyber-insurance markets. In *Workshop on the Economics of Information Security (WEIS)*, Pennsylvania State University, PA, USA, 2014.
- Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy (S&P)*, pages 111–125, Oakland, CA, USA, 2008.
- Roger Nelsen. *An introduction to copulas*. Springer, New York, 1999.
- NetDiligence. NetDiligence 2015 cyber claims study. Technical report, NetDiligence, 2015.
- NetDiligence. NetDiligence 2016 cyber claims study. Technical report, NetDiligence, 2016.
- Hulisi Ögüt, Nirup Memon, and Srinivasan Raghunathan. Cyber insurance and IT security investment: Impact of interdependent risk. In *Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, USA, 2005.
- Hulisi Ögüt, Srinivasan Raghunathan, and Nirup Menon. Cyber security risk management: Public policy implications of correlated risk, imperfect ability to prove loss, and observability of self-protection. *Risk Analysis*, 31(3):497–512, 2011.

- Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. Will cyber-insurance improve network security? A market analysis. In *Conference on Computer Communications (IEEE INFOCOM)*, Toronto, ON, Canada, 2014.
- PCI. Payment Card Industry (PCI) Data Security Standard, v3.1. Technical report, PCI Security Standards Council, 2015.
- John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136, 1964.
- Niels Provos and Thorsten Holz. *Virtual honeypots: From botnet tracking to intrusion detection*. Pearson Education, 2007.
- PwC. Turnaround and transformation in cybersecurity: Key findings from the global state of information security survey 2016. Technical report, PricewaterhouseCoopers (PwC), 2016.
- Sam Ransbotham and Sabyasachi Mitra. Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research*, 20(1):121–139, 2009.
- Sam Ransbotham, Sabyasachi Mitra, and Jon Ramsey. Are markets for vulnerabilities effective? *MIS Quarterly*, 36(1):46–64, 2012.
- Markus Riek, Rainer Böhme, Michael Ciere, Carlos Gañán, and Michel J. G. van Eeten. Estimating the costs of consumer-facing cybercrime: A tailored instrument and representative data for six EU countries. In *Workshop on Economics of Information Security (WEIS)*, University of California, Berkeley, CA, USA, 2016.
- Sasha Romanosky. Comments to the Department of Commerce on incentives to adopt improved cybersecurity practices. Docket number 130206115-3115-01. Technical report, Department of Commerce, 2013.

- Michael Rothschild and Joseph Eugene Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.
- Bruce Schneier. Insurance and the computer industry. *Communications of the ACM*, 44(3):114–115, 2001.
- Galina A. Schwartz and S. Shankar Sastry. Cyber-insurance framework for large scale interdependent networks. In *Proceedings of the International Conference on High Confidence Networked Systems (HiCoNS)*, pages 145–154, Berlin, Germany, 2014.
- Carl Shapiro and Hal R. Varian. *Information rules: A strategic guide to the network economy*. Harvard Business Review Press, 1998.
- Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. Competitive cyber-insurance and Internet security. In Tyler Moore, David Pym, and Christos Ioannidis, editors, *Economics of Information Security and Privacy*, Berlin, Heidelberg, 2010. Springer.
- Gary Stoneburner, Alice Goguen, and Alexis Feringa. Risk management guide for information technology systems. Technical report, National Institute of Standards and Technology (NIST), 2002.
- Xiaomeng Su. An overview of economic approaches to information security management. Technical Report TR-CTIT-06-30, Centre for Telematics and Information Technology University of Twente, Enschede, The Netherlands, 2006.
- Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05):557–570, 2002.

- Kurt Thomas, Danny Huang, David Wang, Elie Bursztein, Chris Grier, Thomas J. Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. Framing dependencies introduced by underground commoditization. In *Workshop on Economics of Information Security (WEIS)*, Delft University of Technology, The Netherlands, 2015.
- Hal Varian. System reliability and free riding. In L. Jean Camp and Stephen Lewis, editors, *Economics of information security*, volume 12 of *Advances in information security*, chapter 1, pages 1–15. Springer, Berlin, Heidelberg, 2002.
- Thomas von Ungern-Sternberg. The limits of competition: Housing insurance in Switzerland. *European Economic Review*, 40(3–5):1111–1121, 1996.
- WEF. Personal data: The emergence of a new asset class. Technical report, World Economic Forum (WEF) in Collaboration with Bain & Company, Inc., 2011.
- Spencer Wheatley, Thomas Maillart, and Didier Sornette. The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B*, 89(1):7, 2016.
- Mingyi Zhao, Jens Grossklags, and Peng Liu. An empirical study of web vulnerability discovery ecosystems. In *Proceedings of the ACM Conference on Computer & Communications Security (ACM CCS)*, pages 1105–1117, Denver, CO, USA, 2015.