

SoK: Cyber Insurance – Technical Challenges and a System Security Roadmap

Savino Dambra
Eurecom

Leyla Bilge
Symantec Research Labs

Davide Balzarotti
Eurecom

Abstract—Cyber attacks have increased in number and complexity in recent years, and companies and organizations have accordingly raised their investments in more robust infrastructure to preserve their data, assets and reputation. However, the full protection against these countless and constantly evolving threats is unattainable by the sole use of preventive measures. Therefore, to handle residual risks and contain business losses in case of an incident, firms are increasingly adopting a cyber insurance as part of their corporate risk management strategy.

As a result, the cyber insurance sector – which offers to transfer the financial risks related to network and computer incidents to a third party – is rapidly growing, with recent claims that already reached a \$100M dollars. However, while other insurance sectors rely on consolidated methodologies to accurately predict risks, the many peculiarities of the cyber domain resulted in carriers to often resort to qualitative approaches based on experts opinions.

This paper looks at past research conducted in the area of cyber insurance and classifies previous studies in four different areas, focused respectively on studying the economical aspects, the mathematical models, the risk management methodologies, and the predictions of cyber events. We then identify, for each insurance phase, a group of practical research problems where security experts can help develop new data-driven methodologies and automated tools to replace the existing qualitative approaches.

I. INTRODUCTION

The modern society is highly dependent on Information and Communication Technologies (ICT). However, despite its paramount importance, the use of ICT also introduces a series of hazards. In fact, computer systems and services are routinely compromised and cyber incidents adversely impact many organizations, hampering business-goal achievements and resulting in copious financial losses [1]. For this reason, cybersecurity has quickly become a subject of debate in executive boards [2] and companies are increasingly investing in ICT security products [3]. Overall, the security sector is expected to grow in 2019 to a 124 billion USD market, with application security testing, data loss prevention, and advanced threat protection representing the core investments [4].

Despite the importance of this considerable and rapidly-increasing effort, it is well understood that cyber attacks cannot be prevented by technical solutions alone and the protection against all possible threats is neither possible nor economically feasible. Thus, in order to handle the residual risk, organizations are rapidly moving towards managing their cyber risk by incorporating cyber insurance into their multi-layer security frameworks. Cyber insurance is defined to be the way to transfer the financial risks related to network and computer incidents to a third party [5]. Compared with traditional insurance policies for business interruption and crime, a cyber-insurance policy can also cover, for

instance, digital data loss, damage and theft, as well as losses due to network outages, computer failures, and website defacements.

A. A booming phenomenon missing solid foundations

As evinced by recent market reports, the adoption of cyber insurance has tremendously increased over the last decade, achieving an annual growth rate of over 30% since 2011 [6]. This is also reflected in the growing number of claims submitted for cyber incidents in a wide range of business sectors [7] and that, in few striking cases, have seen insurance companies paying even hundred-million-dollar indemnities [8].

Following this trend, the cyber-insurance market is forecasted to reach 14 billion USD in gross premiums by 2022 [9] and several indicators confirm this direction. First, cyber crimes have never been so profitable [10] and the growing number of attacks is increasing the awareness of board members about cyber risks and the impossibility of only relying on preventive solutions [11]. This pushes a growing number of companies, among which even more small- and medium-size enterprises, to start considering cybersecurity insurance as a risk mitigation strategy: in fact, data show that 66% of them would need to shut down if hit by a data breach [12]. Another strong driver for the cyber-insurance domain is the introduction of global regulations on personally identifiable information loss, such as GDPR and CCPA. For instance, the need to cover fines and the high cost of handling user notifications are already creating interest in purchasing cyber insurance [13].

In other words, while researchers and security experts are still debating whether cyber insurances even make sense and how they could be better implemented, insurance companies are already selling them as part of their portfolio. We may like it or not, but this is already a reality – and as it often happens in our field, security needs to catch up with an immature technology that was rushed to the market. As we will see in the rest of the paper, companies are currently struggling against the demand of cyber policies as existing tools and methodologies to assess risk exposures and pricing are inadequate in the cyber domain. Although past studies have concluded that, without considering catastrophic scenarios, the vast majority of cyber risks are insurable [14]–[16], carriers are missing solid methodologies, standards, and tools to carry out their measurements. The result, as we will comprehensively detail later in this work, is that purely *qualitative* assessment of such risks leads to inaccurate evaluations, not properly tailored to the customers but mainly based on averages for their industrial sectors [17].

B. Motivation

Researchers and practitioners have studied the main aspects, the evolution, and the core challenges of cyber insurance for more than two decades [18]. Marotta et al. [19] recently published a survey in which they discuss the history, current status, peculiarities, formalization, and future directions of the cyber-insurance domain. However, while researchers have extensively looked at the theoretical aspects of the cyber ecosystem, there exists a very limited number of studies that relied on real data and leveraged the domain expertise of system security experts [20]–[29]. For example, as noted by Allodi et al [30], while in other sectors risk assessment is based on *quantitative* estimations, cybersecurity risks are typically computed by using *qualitative* risk matrices that rely on *subjective* experts opinions. And this is just the tip of the iceberg. Researchers have so far focused on understating if, and to what extent, a cyber-insurance market can be useful, and which advantages and incentives it can bring to the different parties and to the global ecosystem (both in economic and security terms). However, very little has been done to explore how such insurance schemes can be implemented in a rigorous and scientific way.

To cover this gap, this paper aims at providing an extensive discussion of the *technical* aspects and *open challenges* in the cyber-insurance domain, emphasizing how security experts can contribute to this rapidly evolving area. For example, we will discuss how, despite their apparent similarity, risk assessment and risk prediction are not interchangeable concepts, and the method provided by the first may fall short on the requirements for accurately predicting future cyber incidents.

We believe the cyber-insurance field raises many technical questions that require the expertise of system security researchers: how can one identify and collect low-level risk indicators and compare them with externally-observable events? Is it possible to automatically extract dependencies among different software and services and capture the risk introduced by the supply-chain of a company? These are only two examples out of a long list of open research problems we have identified throughout this paper. Our main goal is to present a thorough discussion on these problems such that researchers understand that to work properly cyber insurance will require practical solutions that go well beyond its economic and game-theoretical aspects.

C. Structure of the paper

Our study starts from the description of a classic insurance process for the purpose of identifying its main phases and actors, and clarifying the differences and peculiarities of the cyber domain (section II). We will then look at the existing literature, covering risk management and game theoretical methods, but also economic studies and previous works that tried to *predict* security events. We will try to clearly organize each contribution and point out which part of the cyber-insurance puzzle it tries to address.

In the second half of the paper, we introduce four main research areas where we believe that expertise in computer security can support the cyber-insurance domain. This includes risk prediction (Section V), automated data collection (Section VI), catastrophe modeling (Section VII), and computer forensics (Section VIII).

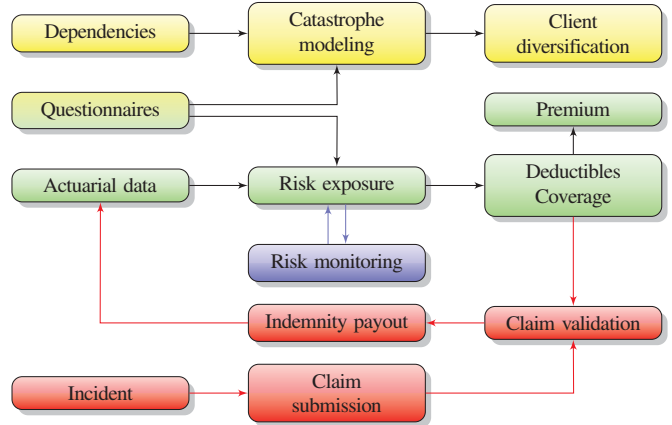


Fig. 1: Classic insurance process workflow extended in a cyber scenario (● Portfolio Management ● Underwriting ● Post binding ● Claiming)

Each section identifies the technical challenges and emphasizes a number of concrete future research directions.

II. PROLOGUE: FROM INSURANCE TO CYBER INSURANCE

A traditional insurance process includes several interacting components, as depicted in the diagram reported in Figure 1. Even though its main concepts and stages might be familiar to most readers, for the sake of completeness, in this section, we briefly provide the basic definitions. This short introduction allows us to then discuss what makes cyber insurance unique compared to all other types of insurances.

The Traditional Insurance process

Insurance is a risk management method whose main purpose is to convert the risk of harmful events into an expenditure. The insurance process generally involves two players: a first supply-side entity who provides insurance, named insurer or insurance company, and a second demand-side entity who buys the insurance, known as insured or policyholder. The two parties interact in two different phases, respectively identified as underwriting (or policy stipulation) and claiming for compensation. During the drafting of a policy, an insurance carrier needs to acquire useful information about the prospective client with the purpose of identifying his risk class. Afterwards, the two parties need to clearly define the conditions, circumstances, and nature of the events that are covered by the policy. Coverage can encompass both first- and third-party losses: while the former is purchased to cover the policyholder against damages or losses suffered by the insured to his person or property (e.g., health, disability insurance), the latter is intended to protect the policyholder against liability for damages or losses caused by the insured to other people or their property (e.g., bystanders hit by insured’s car in an accident, stranger’s properties damaged by a fire that comes out of insured’s house). At this point, the insurer quantifies the material damage that the insured — or third subjects if considered — would be subjected to if these occurrences were to happen. Finally, the insurance company takes on the liability and management of such situations cashing a premium payout from the insured.

The management of client portfolios is another crucial task insurance companies need to consider during the underwriting phase. The goal is typically to maintain a pool of policies, each of them having an independent probability of claim. This diversification averts catastrophic scenarios in which a single incident impacts a large fraction of the clients: in such cases, a significant number of claims would be submitted at the same time and the insurance would suffer a huge blow in covering losses. For instance, it may not be a good strategy for an insurance company to insure against fire hazards all apartments located in the same building.

Finally, when experiencing losses due to an incident which is potentially covered by the insurance policy, the victim submits a claim to the insurer who makes sure of its validity, assesses the impact of the event and compensates the claimer with an indemnity determined according to the terms of the policy. The contract can also include a deductible, i.e., an amount for which the insured is liable on each loss.

In order to make this entire process possible, the insurer must carefully set its tariffs to ensure that the premiums collected are enough to cover future claims, in addition to yield profit for the insurance firm itself. Unfortunately, this is anything but easy. Indeed, when it comes to selling a finished product or service, a firm can easily determine its price knowing which costs have been incurred for its realization. On the contrary, an insurer who places its product on the market does not know in advance the amount of money required for claim compensations because of their inherent uncertain nature. In this respect, actuarial techniques allow to estimate these disbursements and overcome the cost uncertainty related to this inverted production cycle. A key element for this estimation relies on statistical methods that study how claims for covered events have evolved over the previous years to forecast their future evolution. Thus, the raw information required to build a classic insurance product consists of a large set of historical records containing claims and compensations for events which have similar characteristics to the ones being insured. Insurance firms usually do not rely only on their own data sources but also take advantage of the market statistics that aggregate historical data of other companies in the same domain. This statistical information, which normally goes under the name of *actuarial data*, is what allows an insurance company to estimate the risk of a certain event or client, given a number of relevant contextual information (acquired during the underwriting phase). This includes, for instance, the driver's age and neighborhood for a car insurance or the age of the building in a house insurance. Unfortunately, as we will discuss next, actuarial data are scarce in the cyber domain and the characteristics that need to be collected about a client (and that presumably are good predictors for future incidents) are not yet well understood.

Extending Insurances to the Cyber Domain

With the help of Figure 1, we now look more closely at how the previous process is applied to the cyber domain by discussing the differences and the main challenges that affect each insurance phase.

Underwriting – As we discussed above, the policy underwriting requires the insurer to collect information from the client that can be useful for the purpose of risk assessment. Following a traditional model, also in the cyber domain this is still performed by a mix of

self-assessment questionnaires, checklists, business documentation, meetings, and interviews [31]–[36], whose objective is to identify the adopted software and technologies, the deployed security measures, the presence of sensitive data and how it is stored and processed, and any other information that can affect the global security posture of the company under investigation [37], [38]. A deeper analysis can be carried out to tailor the product to the specific customer based on its characteristics and requirements: a monitoring software equipment together with an overhaul of preexisting security logs and telemetry serve this purpose. Finally, some deficiencies and precautions are often advised to the client in order to comply with the best-known security practices [39].

Assessing the cyber risk of organizations or individuals is an overly challenging problem due to a number of reasons including the existence of asymmetric information, the dynamic nature of the cyber ecosystem, and the indirect risk that might be propagated from the relations with the third parties (we will come back to these points in more details later in the paper). Although with the traditional meticulous risk assessment methodologies the underwriters could draw an approximate picture of the customer's risk exposure, they might not be aware of the residual risks that might be known to the counterparts. The possession of a greater material knowledge by one of two parties involved in an economic transaction creates the problem of carrying asymmetric information and this represents a major issue in cyber insurance [5], [16], [40]–[42]. A risk assessment that is made by analyzing asymmetric information can lead to adverse selection [43], [44]. For example, unfair risk scores might be assigned to a company whose private and inaccessible information may reveal a severe exposure to risk compared to another with a better security hygiene.

The existence of asymmetric information also impacts negatively the customer side as insurance firms may raise premium prices due to incomplete knowledge and risk overestimation, leading to an expensive, niche, and not-appealing product [45], [46]. High premiums are also the result of insufficient criteria to reduce them: even if a company holds security certifications and profusely invests in self-protection, the effectiveness of these actions against the wide variety of cyber attacks is not clear, making, in turn, difficult to assess to what extent they are useful to reduce the overall risk [45]. A timid step in this direction is the one of some carriers who reduce premiums or deductibles if the client uses risk assessment tools, security technologies, and breach response services of specific vendors [47].

The interdependent nature of the cyber ecosystem makes the risk estimation even more complicated. Nowadays, when cloud computing and outsourcing are two mainstream phenomena, cyber risk is intertwined among all entities that depend on one another [5], [41], [48]. Companies may indirectly get damaged because they use external services that are targeted by a cyber attack: an example is the recent DDoS attack against DynDNS – which impacted more than sixty of its customers [49]. Thus, a firm's measures and expenditures in self-protection may not proportionally increase its security level when making use of services from third parties that do not invest as well [5], [18]. In the pre-binding phase, risk exposure must be then identified from a holistic standpoint, preferring a due diligence approach to a simple checklist and including in the review all internal and external threat vectors that could potentially

compromise pre-insured's security [50].

Actuarial and Pricing – The actuarial approach based on statistical models described above does not fit the cyber domain where historical data of claims and compensations are still scant [16], [41], [43], [45], [51], [52]. Enterprises experiencing a cyber incident have a strong incentive not to publicly disclose it as this would tarnish their image. As a result, the few available databases [53]–[55] contain records which are often vague, missing details, and biased towards large and serious incidents, whose disclosure is unavoidable due to their resonance or due to mandatory-notification laws [56], [57]. The infeasibility of the actuarial approach alone for an accurate risk estimation is corroborated by its ever-evolving components: cyber threats and attack methods swiftly evolve alike defense methods and strategies do [18], [38], [41].

Portfolio Management – As briefly discussed before, a fundamental requirement of traditional insurance schemes is that the insurer should strive to obtain a portfolio of policies with an independent probability of claim submission. This diversification can reduce the likelihood that a single incident could harm a considerable portion of clients – a *catastrophic event* that can have severe consequences and cause the bankruptcy of the insurer [58]–[60]. Unfortunately, it is harder to obtain such diversified portfolio in the cyber domain, due to the monoculture of software and hardware products [61]–[64]. Although deploying different configurations is possible, recent events have shown that the business continuity of a large set of possible clients – independently of their size, sector, and assets to protect – is simultaneously undermined when a piece of a broadly-used software or hardware is found to suffer from a severe vulnerability [65]–[72].

In other domains, a common way insurers protect themselves against catastrophic events such as wildfire and hurricanes is by purchasing policies from other insurance companies. Sadly, the current lack of re-insurers in the cyber domain further exacerbates this problem [41], [73]–[75].

Post-Binding Phase – Due to the complications in both the policy underwriting and claiming phases, an additional post-binding phase is introduced, which does not exist in other forms of insurance [50]. In fact, in traditional insurances, the relationship between the firm providing coverage and the policyholder ends once the contract has been signed and the two parties interact again only in case of a claim submission. On the contrary, a cyber insurance may require periodic risk assessment after the underwriting is completed, to allow the insurer and the policyholder to collect updated information related to new threats and evolved risks. Indeed, many cyber-insurance policies already bring supplemental value through the inclusion of risk mitigation, tracking and loss-prevention tools [76]. Clients, in particular small organizations that lack experience, can benefit from this continuous interaction to better ponder their measures towards higher-priority situations [46]. The post-binding phase also helps to prevent the well-known issue of moral hazard [18], [43], [77], [78] — a form of post-underwriting opportunism by the policyholder, who undertakes incautious actions knowing that, in case of incidents, there exists a counterpart who will bear the brunt and will not be able to verify the presence of negligent and fraudulent actions. In this regard, insurers have to conduct continuous risk assessments to

resize the set of inaccessible information of the insured and mitigate its unfair behaviors.

Claim Submission and Validation – Cyber-insurance policies usually cover the costs of incident response and forensic investigations, including the identification of stolen or compromised data and the extent to which third parties have to be informed according to the current regulations. Despite this, a precise quantification of the involved and compromised assets is complicated by their intangible nature [41], [43]. In addition, since jurisdictions may apply different notification laws, each case must be accurately evaluated according to the localization of the indirectly-damaged third party.

The insurer as well compensates for economic losses related to the event. In particular, cyber insurance may refund losses due to business interruption caused by an attack, as well as cyber extortion and stolen assets. This approach is insufficient in the cyber scenario where the above primary losses are often followed by secondary ones that result from a loss of reputation whenever the incident is publicly disclosed [40].

Time is also a key component when it comes to claim submission. Some attacks may silently compromise a system and remain undiscovered for a considerable time-frame. The validity of claims in such situations is a more arduous issue to formalize in cyber policies. Furthermore, carriers may require forensic investigations prior to claim submission to verify its validity, resulting in an initial disbursement from the insured and a reputation damage due to the disclosed incident.

III. LITERATURE ON CYBER INSURANCE

A. Categorization and Source Selection

Since its first appearance in the late 90s [18], cyber insurance has been the focus of researchers from different disciplines.

For our study, we selected and analyzed 93 works among academic papers, standards, and frameworks. As shown in Figure 2, we grouped these works in four main categories and fourteen sub-categories. In particular, we found that previous research has mainly focused on two areas: cyber risk management, which tries to estimate attack probabilities and possible damages, and mathematical modeling and game theory simulations, which aim at deriving interesting properties on the consequences of cyber-insurance adoption. Two additional areas complete the picture: research conducted by the economics community reporting figures from past incidents or discussing the costs of possible scenarios, and research focusing on the prediction of future cybersecurity events.

Since these four macro categories refer to very different research domains, we adopted distinct criteria to select and present the contributions from each of them. *Risk management* is a very wide topic that covers a wide range of domains, ranging from pharmaceutical products to natural disasters. We reported all methodologies and frameworks that are currently used in IT, together with those academic papers presenting risk aggregation techniques. Regarding the contribution from the *economics* community, as an exhaustive discussion would be out of scope for a security conference, we focused on the papers needed to emphasize research problems, existing tools, and on the major findings that can affect the work of security researchers. For this reason, we comprehensively reported




Fig. 2: Cyber-insurance research areas

all of the attempts made in quantifying economic losses following cyber incidents. The works employing *mathematical modeling and game theory* have already been deeply analyzed by Marotta et al. in [19]. Therefore, in Section III-C, we decided to offer a different systematization that focuses on which property the authors were interested to prove, along with the choice of the simulation parameters —e.g., the market model, the presence of asymmetric information, and the network topology. Finally, since our ultimate goal is a call to actions for security researchers to provide data-driven solutions for the cyber-insurance domain, our study comprehensively presents and compares previous *prediction* attempts in section III-E.

B. Approaches and techniques for cyber risk management

According to the ISO standard 31000, a risk management process can be described as a set of tasks whereby it is possible to measure the risk and subsequently develop strategies to monitor and control its evolution [132]. As a result, the first phase of risk management is dedicated to the identification of the valuable assets and of the related threats that represent the main components of risk. Each threat is then analyzed by evaluating its likelihood and possible impact from both a qualitative and quantitative perspective, and results are then aggregated to obtain an overview of the whole risk. These two phases, grouped and referred to as *risk assessment*, are usually followed by a risk treatment step, which covers the choice of non-exclusive countermeasures that can be adopted to tackle each of the risk components. Finally, as risks may suddenly change, causing the previous estimations to become incorrect and countermeasure ineffective, a number of monitoring and reviewing actions are required to continuously update the risk estimation.

Risk management is an important process when it comes to information technologies. Therefore, the literature is rich of guidelines, frameworks, and techniques that contextualize it to the digital world. As depicted in Figure 2, we grouped under the risk management sub-category the studies that provide a walk-through of the entire procedure [32]–[34], [95]–[98], defining terms and providing a helpful documentation of how to address issues on risk assessment and treatment, as well as insights on risk monitoring and reviewing. Other works often inherit or revisit a previous risk management methodology and introduce new techniques to implement a specific sub-component. In this respect, we created two different sub-categories in which we respectively list the works addressing the whole risk assessment [36], [99]–[101] and those narrowing down the discussion on risk analysis [35], [102], [103].

Although widely used standards (such as ISO 27005 [95]) and tools (e.g., NIST SP 800-30 [32], Magerit [33], OCTAVE Allegro [34], Clusif [36] and the one proposed by Microsoft [97]) handle the single stages of the management procedure in a different way, they share a common underlying workflow for assessing individual risks. According to it, the process typically starts by brainstorming which and how cyber-based threats could prevent the company from reaching business goals and team objectives. In this respect, real-life cyber events previously occurred to other companies can be used as source of inspiration. In addition, frameworks often provide guidelines on how to identify this collection, including checklists or questionnaires, and advising to adopt a *what-if* approach to understand what could go wrong and what the possible consequences are. The outcome of this process is the

creation of a *risk register*, whose structure, together with some examples, is reported in Table I. Once each row has been filled with a description of the threat, including its possible triggers and effects, the impact and likelihood of its materialization are assessed to define the *inherent* risk. Two approaches exist for scoring these factors and the choice of one rather than the other depends on the company itself. Indeed, some tools provide a table of decipherable words with a *qualitative* description, whilst others opt for a *quantitative* numerical sliding scale (e.g., Table I3 and H3 of [101]). It is worth pointing out that the same event could be assigned different values across distinct situations: if an organization’s public statement is “*we have built our reputation on our commitment [...] to protect the privacy and confidentiality of personal information*”, the impact of user-data leaks for this company will be higher if compared to another one with different prerogatives. The next step is the identification and mapping of existing mitigations or controls that could reduce the likelihood of each threat: companies often take advantage of existing frameworks that list critical checks and best practices, and indicate the extent to which the control environment reduces the inherent risk. As a result, a value reflecting the *residual* risk is obtained and a three-fold choice opens up: if the value falls within the company’s risk appetite limit, no further action is needed in this phase. If not, more controls and mitigations have to be investigated or the residual risk has to be transferred to a third party —e.g., with a cyber-insurance policy.

Finally, in the last sub-category of Figure 2, we reported all the methodologies that have been proposed to *aggregate* and propagate individual risks based on tools that capture the relationships among different information components or requirements of an attack. These modeling tools make use of graph theory or model checking to draw conclusions starting from some preconditions. Among them, attack trees are widely used techniques to capture dependencies among threats [104]–[109]. Each tree is a leveled diagram made of nodes, leaves and a root; each node represents an attack or a threat which materializes only if all its children are satisfied. The root attack is completed if all nodes are satisfied. Similarly, vulnerabilities or exploits are represented as nodes in attack graphs and conditionally linked to each other according to their preconditions and results. Such composition of vulnerabilities is used to simulate incremental network penetration and attack likelihood propagation with the purpose of measuring the overall security of a system or network [110]–[113]. Finally, hazard and operability studies (HAZOP) [100] and failure mode and effects analysis (FMEA) [114], are other two techniques used to break down a complex process into small sections and reason about possible undesired situations, their causes and consequences. Such kind of tools are mostly employed when the use of ICT can introduce a series of hazards in industrial environments [115]–[117].

As we will discuss later in this section, these methodologies, inherited from other domains, can be unsuitable when employed in cyber scenarios.

C. Cyber insurance and Game Theory

A large portion of existing contributions employ mathematical modeling and game theory to infer properties and effects of adopting cyber insurance. As comprehensively reported in [19], this approach allows in the first place to create a mathematical model of cyber

insurance which takes into account its main actors (insurance carriers, policyholders and regulatory entities), their interdependences (probability of infection and externalities), the network topology (independent nodes, complete graph, random graph, or others) and the market type (competitive, monopolistic, or oligopoly). Once a model has been defined, game theory is used to simulate the behavior of agents: insureds choose their desired level of protection and contract type, insurers instantiate contracts, and regulators come into play by imposing regulation options (mandatory insurance, fines, bonuses, penalties, mandatory investment, etc.). The use of game theory makes it possible to also include in the models the major issue of information asymmetry in its moral hazard and adverse selection forms. This way of tackling cyber insurance is very useful for strategic purposes and allowed researchers, practitioners, and governments to reason about consequences and peculiarities of its employment, and market viability.

1) *Viability of the cyber-insurance market*: As already discussed, the starting point of each simulation is the definition of a mathematical model of cyber insurance that considers its main aspects, e.g. market type, type of coverage, existence of asymmetric information, network topology, etc. Therefore, an important finding of each simulation is to verify whether the market defined by such pre-conditions may exist or not, i.e., whether the actors would opt for the insurance case over the non-insurance one. One way to achieve this result is the comparison between the average utility function for agents with ($E[U^I]$) and without ($E[U^N]$) insurance: in the economic theory, this function measures the welfare or satisfaction of an entity from consuming a certain number of goods. Then, if $E[U^I] \geq E[U^N]$ holds, the choice of an insurance policy directly contributes to increase the wealth of an agent [19]. Almost all previous works —among which we find the more realistic settings that include a competitive insurance market, non-zero-profit carriers, the presence of moral hazard and adverse selection, and a partial coverage whose level is defined by the policyholder— fall in this category [42], [80]–[82], [84], [86]–[89], [91], [93]. Only two studies found that actors who decide not to invest in a cyber policy would benefit from this choice [93], [94]. Yang and Lui [93] concluded that cyber insurance is not a good incentive for all nodes when modeling a competitive market with zero-profit carriers only offering full coverage and accepting asymmetric information in its moral-hazard form. Naghizadeh and Liu [94] simulated instead a monopolistic profit-neutral insurer, acting as a regulator that imposes fines and grants rebates, and found that this leads to a market failure because of agents would not voluntarily purchase any insurance.

2) *Consequences of cyber-insurance employment*: Among the main topics of interest in this area, we find the use of cyber insurance as an incentive for internet security [79], [83], [84], [86]–[90], the change in self-protection investments when insurance is available [79]–[85] and its contribution for reaching the social welfare [79], [87], [91], [92]. These studies concluded that cyber insurance is not a good incentive for internet security in presence of a competitive or monopolistic market and asymmetric information in its moral hazard form [42], [83], [84], [86], [87]. On the other hand, researchers also concluded that a non-competitive cyber-insurance market can increase internet security if fines are imposed by regulation entities and policy are carefully designed.

TABLE I: Risk register: qualitative assessment examples for inherent and residual risk

Description	Cause	Effect	Inherent Impact	Inherent Likelihood	Inherent Risk	Residual Impact	Residual Likelihood	Residual Risk
Third person gains access to sensitive customer information via stolen credentials	Employee inadvertently inputs access credentials within the source code	1 million customers at risk of identity theft Company receives significant criticism for its privacy preserving policy	Catastrophic	Possible	High	Catastrophic	Remote	Medium
Sensitive customer data exposed to unauthorised parties	Employee deliberately copied full customers records motivated by personal financial gain	1 million customers at risk of financial theft	Catastrophic	Remote	Medium	Catastrophic	Extremely Remote	Low
Remote code execution on webserver by unauthorised parties	Zero-day vulnerability exploited in third-party library used for customer authentication	1 million customer data at risk of theft Online platform not available to customers Business-continuity interruption	Catastrophic	Possible	High	Catastrophic	Possible	High

When analyzing the effect of employing cyber insurance on self-protection, some works show that, if insurance is available, agents prefer not to invest in self-protection, but rather in insurance contracts [133], [134]. In this case, minimal investments imposed by regulators do not change the results. Finally, the usefulness of insurance as a tool to reach the social welfare and the optimal level of self-protection investments has not been yet understood: different studies [88], [135] reached contradictory conclusions on this topic although considering the same preconditions, probably because of adopting different network topologies—which lead to different interactions among actors—throughout their simulations.

D. The Economics Perspective

Since cyber attacks are often considered inevitable events, cyber experts are increasingly focusing on their economic consequences [119]. In this respect, scenario-based evaluations are a very common approach used to serve two main purposes. For a company, these scenarios provide a useful way to assess the possible consequences of a cyber event [123], to measure the incident response capabilities [122], and to identify the critical systems, people and premises that are needed to continue to serve their customers [120]. On the insurance carriers side, simulations based on scenarios are often used to estimate the financial impact of large-scale attacks or catastrophic events that hit many businesses at once [121]. This simulation practice is rapidly gaining popularity due to the current cyber landscape, in which the costs of recovering from particular types of attacks are way greater than the cost required to prevent them [118]. Furthermore, tests can help companies to emphasize the presence of valuable data to protect and shed light on interconnected risks that could lead to catastrophic events [119]. Good evidence of this can be found in the decision of the European Insurance and Occupational Pensions Authority (EIOPA) to include, for the first time in 2018, cyber scenarios in the collective insurance stress test used to assess cyber-risk response [136].

The creation process of a scenario-based simulation goes through a multi-stage procedure [60] and it is usually performed by C-Suite executives due to their expertise in business-critical roles and operations [119]. The process starts with the creation of a plausible scenario, defined by a footprint of events to be simulated and a contagion mechanisms among the involved entities [124]. There is a wide range in the type of scenarios that can be used for different

applications. For instance, scenarios can be based on historical or synthetic events, they can be generic or specific for a given company or sector, and they can consider single or multiple events [125]. Scenarios allow the simulation of both common digital incidents—like data exfiltration, cyber extortion, denial of service attacks, financial transaction compromise, and cloud service provider failure—as well as rare events—such as cyber-induced fires in buildings or industrial plants, cyber theft of marine cargo, cyber attacks on power grids, or oil rig explosions due to platform control system (PCS) compromise [126].

Developing a scenario is a challenging task as it is not easy to fully understand all the systems involved and predict the possible cascading effects that could be triggered [60]. For this reason, developing a *coherent* scenario is a key aspect for successfully achieving the second phase of the simulation that consists of estimating the induced losses to a business or the impact of claims submitted to an insurance company by taking into account its client portfolio.

The output of the simulation can be further extended beyond a single company by taking into account macroeconomic consequences too [124]. This result can be achieved by selecting a representative subset of the whole population of companies from a wide range of different business sectors and use them to estimate the losses of a given scenario. In turn, this allows for a quantification of the effects on many variables of the global economy [60].

Besides scenario-based simulations, other economic studies attempted to gain insights into cyber risks by leveraging publicly available data. For instance, Eling and Loperfido [128] analyze statistical properties of a data breach information database to show that data breaches significantly differ among each other, hence they must not be put in the same basket but must be mapped to separate risk categories. Using another dataset of publicly available survey data, Herath et al. attempted instead to build a pricing model for cyber-insurance premiums with the robust copula methodology [127]. Premiums for first-party losses due to virus intrusions are computed with a probabilistic model based on three factors: the occurrence of the events covered by the policy, the time from the issue of the policy to the incident, and the indemnity paid by the insurance in case of the breach occurring. Biener et al. [16] analyzed the world’s largest collection of publicly reported operational losses to draw empirical conclusions on whether cyber risks are insurable

or not based on Berliner’s criteria. Results suggest that cyber risk owns some peculiarities that undermine its insurability, such as its evolving nature, the lack of actuarial data and reinsurance, the severe information asymmetries, the limited coverage and caps, and the high deductibles and premiums for small and medium enterprises.

Wheatley et al. [129] statistically modeled a 15-year cyber-breach dataset to show that the size of an organization is strongly coupled with the frequency and severity of breaches, and the number of information leaked during such events is expected to double within five years from two to four billion items. The handling and response costs of two data breach events are at the center of the study by Layton et al. [130]. Counterintuitively, the authors show that none of the two incidents negatively affected the company stock price and economic growth, secondary and intangible losses have negligible importance with respect to direct losses, and policy and procedure for handling the event have a large effect on the overall cost. On the contrary, in [131] security breaches are found to negatively impacting stock quotation of the victims, especially in the case of e-commerce firms and DoS attacks.

E. From risk assessment to risk prediction

So far, a considerable amount of studies, frameworks, and methodologies have focused on assessing the risk of cyber attacks by explicitly defining their underlying causes and triggers. In fact, as we show in Table I, the first column of each row specifies either the particular action, the vulnerability or the exploit that makes the risk materialize. While this *assessment* technique is well established in other domains (e.g., industrial and financial), its effectiveness is still unclear in a cyber scenario. Indeed, if the whole evaluation is based on the *current* knowledge of vulnerabilities present in the system and tools, and on the exploits available to the attackers, it quickly becomes clear that the final measurement has limited lifespan, as new ones are respectively discovered and released on a daily basis. Moreover, when major cyber incidents occur, its root causes and enabling factors are almost always *unknown* to the community, greatly complicating the assessment of the associated risk.

The goal of *prediction* is to overcome this assumption and carry out the risk estimation by leveraging a combination of *risk indicators*, measurable factors that have been empirically proven to reflect the risk across a number of experiments. For instance, back to Table I, lower age, frequent use of untrusted internet connections, and longer browsing sessions at night have been found to be good signs for predicting which users are more at risk of malware infections [26]. And this is done by mentioning none of their actions or incautious behaviors — e.g., the user clicked on a malicious banner or installed malicious software. In a similar way, companies with misconfigured DNS services and expired certificates more frequently show signs of botnet activities, otherwise less likely to be observed in other entities where those are correctly set up [20].

These measurable indicators are merely correlated and not the cause itself of the risk, the same way as the driver age is not the cause of car accidents. But by measuring these signs, experts can make predictions of the likelihood of future events.

Over the past two decades, few scattered studies have focused explicitly on the problem of predicting security-related events. In 2001, Browne et al. proposed a simple formula to predict the

amount of security incidents, as a function of time, related to a known vulnerability [23]. Bozorgi [24] used instead publicly available vulnerability databases to predict which, and how soon, a vulnerability is likely to be exploited in the future. In 2005, Schechter [137] looked at the challenges of predicting cyber attacks. He discovers that experts had a much better understanding and success in modeling traditional crimes, such as home burglary [138] while “*attempts to bring the quantitative approaches of insurance and risk management to the measurement of [computer] security risk have failed*”. The author concluded that this is due to the fact that we still lack techniques to measure the security strength of a piece of software (we will get back to this idea of predicting risk through measuring security in Section V).

Another traditional way to predict future events is to adapt *software reliability growth models* (SRG) commonly used by the reliability community to describe (typically through a non-homogeneous Poisson process) and predict the evolution of defects in a software artifact. For instance, Condon et al. [139] show that specific classes of computer incidents (such as those that depend on particular vulnerabilities) can be modeled with an SRG, while the total aggregated incident rate can be better approximated by using time series [140].

In 2016, Edwards [141] found that the daily frequency of data breaches can be described by using a negative binomial distribution and used this model to estimate the likelihood of similar incidents in the future. Maillart [142] found instead that the theft of personal information follows a power-tail distribution that is robust independently of the sector and size of the targeted organization.

On a different but related topic, a large corpus of works aimed at predicting the occurrence of new vulnerabilities in software products [143]–[148]. However, as we will discuss in Section V, it is still unclear how this information can translate to a prediction or the likelihood of being attacked or compromised in the future.

In recent years, prediction techniques have been at the center of few works for the purpose of assessing the risk in different circumstances. In 2009, Bossler et al. [29] investigated the influence of different factors in predicting data losses from malware infections by conducting a survey over 788 college students. More recently, Liu et al. [20], by using a set of external observable features, attempted to predict the likelihood of an organization to suffer a cyber incident in the future. The authors achieved, overall, a 90% accuracy with 10% of false predictions. Cyber incidents are considered also by Thonnard et al. [22], who discussed organization- and individual-level features that are correlated with the risk of experiencing spearphishing attacks. In a similar way, Sarabi et al. [21] build a predictor for cyber incidents using a set of industry, business, and web visibility/population information. RiskTeller [25] is a prediction tool that leverages internal telemetry data to predict which machines are at risk of being infected by a broad spectrum of different malware. Its prediction accuracy reaches 95%, showing that such tool could be used to prioritize security spending towards machines at higher risk of infection. The same conclusion is reached by Yen et al. [26], who use logs from an antivirus software to infer the risk for hosts in a large enterprise to encounter malware. On the consumer side, Canali et al. [27], assess to what extent the risk class of a given user can be predicted based only on his web browsing behavior. The authors show how certain types of user actions

TABLE II: Works on prediction

Year & Paper	Predicted event	Ground truth		Features	Feature datasets	
2015 [20]	Cyber incidents	Incident reports	Ext	Mismanagement signs Malicious activities	Scanning tools Public scan data	Ext
2015 [21]	Cyber incidents	Incident reports	Ext	Website statistics Industry sector Size Region Popularity	Information services	Ext
2001 [23]	Vulnerability incidents	Incident reports	Ext	Exploit release timing	Vulnerability database	Ext
2010 [24]	Vulnerability exploitation	Vulnerability reports	Ext	Vulnerability features	Vulnerability reports	Ext
2015 [22]	Targeted attacks	Mail scanning service	Int	Industry sector Size Employees features	Industry classification Linkedin Int telemetry	Int + Ext
2017 [25]	Malware encounters	AV Telemetry	Int	Binary file appearance	Int telemetry	Int
2014 [26]	Malware encounters	AV Telemetry	Int	Demographic VPN logs Network logs	Int telemetry	Int
2007 [27]	Malicious website encounters	AV Telemetry	Int	Browsing behaviors	Int anti-virus service	Int
2018 [28]	Malicious website encounters	Website Blacklist	Ext	Browsing behaviors Self-reported data	Mobile ISP tracking data User questionnaires	Int
2009 [29]	Losses from malware infection	User questionnaires	Int	Routine Activities Deviant Behavior Guardianship	User questionnaires	Int

considerably affect their risk exposure. In a similar way, Sharif et al. [28] use mobile users' browsing patterns complemented with self-reported data to predict whether the users will encounter malicious pages on a long and short period of time. In the latter case, on-the-fly predictions within a browsing session could be useful to proactively prevent malicious-content exposures. All these prediction efforts are summarized in Table II, alongside the type of predicted events, the source of ground truth information, the adopted features, and the data from which they are extracted. The table also shows if the ground truth and the predictive features are extracted from internal sensors (Int) or are measured from public external information (Ext). We will return on the importance of this aspect in Section VI.

Finally, few studies have focused on predicting the cost of cyber incidents and data breaches. In this area, Jacobs [149] proposed a regression model based solely on the number of user records compromised. Romanosky [150] introduced more variables (including the revenue and company type) and found that a 10% increase in firm revenues is correlated with a 1.3% increase in the cost of an incident. The author also noted that the price is ultimately related to the size of the company and the size of the breach, and not to the malicious nature of the incident or its outcome.

F. Discussion

Nowadays, cyber risk management methodologies, results of game theoretical studies, and scenario-based simulations are key components for the development of the cyber-insurance market. In the first case, companies and individuals that want to adopt cyber insurance can take advantage from the existence of these frameworks and guidelines, despite the fact that they were not designed with the insurance market as ultimate goal: indeed, risk management plays a very important role to estimate attack probabilities and possible damages, allowing, in turn, individuals and companies to reason on their needs for a cyber policy. Insurance carriers as well use these tools during contract underwriting for assigning a value to a certain entity's risk and compute premiums for cyber-insurance policies [37].

Unfortunately, all available solutions discussed above have a *qualitative* foundation and base their analysis, assessments, and consequently their results on metrics based on experts knowledge and previous experience, missing a feedback from real-world experiments and measurable quantities. Existing methodologies rely on checklists, worksheets, knowledge basis, catalogues, tables, and what-if reasoning for identifying threats and hazards. The value of this type of analysis largely depends upon the quality of the used documents and the experience of the experts who brainstorm about undesired events and their effects. In the same way, the use of tools to capture dependencies among threats such as fault trees or the outcome of HAZOP and FMEA studies also assumes that who carries the analysis has detailed knowledge about the areas, operations, and processes that may be exposed to hazardous events and conditions.

The absence of objective measures and the qualitative nature of these methodologies make it also harder to obtain an actual value for the likelihood of a given threat in a cyber scenario: threat probability is, in fact, a key component for assessing risks and, although simulations can approximate the frequency of popular attacks found in the wild, the limitations discussed in the actuarial paragraph of Section II exacerbate the quantification of such quantity.

Finally, since a sheer number of risk assessment methodologies exists, it is still unclear which one fits best the cyber domain and provides the most precise way to compute the likelihood of cyber incidents. This aspect is further exacerbated by the ever-growing adoption of IoT devices, for which new risk metrics and specific risk evaluation methods are still missing [151]. Very similar considerations apply to simulations based on scenarios, as their creation, refinement, and precision to capture the intricate relationships among different entities depends completely on qualitative opinions of expert users and C-suite members.

As risk management methodologies and scenario-based tests, game theory applied to cyber insurance can provide important practical insights. Nevertheless, all conclusions obtained from these studies are purely based on mathematical modeling, with all the

limitations that this implies. First of all, the finiteness of modeling can lead to a huge difference between the actors considered and their actual number. Moreover, when using game theory to simulate the behaviors of clients and carriers, players can undertake a limited set of actions and interact with each other only in pre-defined ways, defined by assumptions respectively on the market type and network topology. Unfortunately, there is no measurement or comparison with real-world data that confirms the validity of models and veracity of game theory results.

IV. FROM THEORY TO PRACTICE

As we described in the previous section, research in the cyber-insurance domain has mostly focused on theoretical studies (from a mathematical viewpoint) and on the analysis of the costs/benefits tradeoff (from an economics viewpoint). At the same time, the system security community has instead been largely ignoring this emerging area. This could be simply the consequence of the lack of interesting problems that require novel and practical solutions, or it could be due to the lack of awareness from our community towards these problems. As we believe the latter to be true, we now focus on some of the areas where researchers' experience with system and network security can play a fundamental role to help the development of the cyber-insurance domain. The contribution of system security researchers can help the development of quantitative, data-driven methodologies, and it can bring automation and support tools to replace questionnaires and qualitative estimations.

In particular, we selected four classes of problems, one for each of the insurance phases: actuarial, underwriting, portfolio management, and claim validation. For each class, we underline the limitations in the current approaches, discuss the challenges of proposing new solutions, and outline a number of open research directions for researchers in the security field. To ease their identification, we tried to mark the main open problems we discuss in the text as $\langle \mathbf{R}_i \rangle$.

V. AREA 1: RISK PREDICTION

“They could tell you exactly the chance of an office building burning down in Midtown Manhattan, but there isn't anyone on this planet who could tell you the probability of a large U.S. retailer being hacked tomorrow”

– Graeme Newman, Director at CFC Underwriting [152]

Cyber-insurance providers employ underwriting tools to collect the information required to differentiate the risk across all the applicants [37]. Today, underwriting questionnaires ask a number of questions which insurance companies believe to be relevant to classify the risk of a potential customer. However, as we discussed in Section III-E, researchers still have to identify reproducible ways to estimate risk based on a number of observable features that had been proven to be meaningful predictors across a number of experiments. The experiments conducted to date were often inconclusive and difficult to compare as they were all conducted on different datasets and none of them was ever repeated or validated by other studies. As a result, as a community we still lack an understanding of *which security events can even be predicted in the first place, and which features are most useful for such prediction*. This opens several research directions to explore different methodologies to capture and aggregate risk factors.

Measure the security posture of the target. One of the first ideas that comes to mind to understand the risk of cyber incidents is to look at the overall security of a given target. In fact, the security posture of an organization may provide good insights on the level of risk – *if* we assume that a better security hygiene can lower the risk of future attacks. Indeed, at least intuitively, the higher is the security of a system, the lower should be the probability of a security incident affecting that system. If we accept this assumption, risk prediction can be re-formulated as a problem of *measuring security*.

While the fact that security countermeasures could result in a reduced amount of computer abuse was first assessed in 1990 by the seminal work of Straub et al. [153], the link between security posture and cyber risk is not so straightforward and it is still poorly understood today. Security measures can certainly raise the bar for the attackers, but risk also depends on the number of attacks a target may receive—which could be higher for large and popular organizations. Moreover, relevant targets may attract more sophisticated and motivated adversaries, which can make prediction more complicated. But even if we accept this premise to be correct, there are still two serious obstacles to this approach.

First, despite almost four decades of attempts, it is still unclear whether a way to *quantify* security even exists [154]. For instance, in 2009 Verende et al. [154] surveyed many techniques taken from the economics, the computer science, and the reliability community, but still found unclear the validity of the existing results. Second, even if we had a scale to precisely measure security, it is still unknown what is the exact relationship between the level of security and the probability of incidents ($\mathbf{R1}$). Simply saying that more security equals less risk is too vague to be practical. Does doubling security reduces the risk by half or by a factor of four? Does the curve reach a plateau, after which adding more security does not provide a tangible reduction in terms of risk?

Measure the behavior of the target. The fact that the *behavior* of the target can considerably affect its overall risk is another aspect which is often taken for granted. The idea is that, regardless of its security posture, the risk of being compromised of a given entity increases simply because of the actions it performs. For instance, if a user spends a considerable amount of her time browsing dubious and less reputable web sites, it seems reasonable that she would incur higher chances of being infected by malware than a user who only browses corporate and popular sites. Unfortunately, even if this may seem a logical conclusion, researchers have struggled to measure this simple relationship ($\mathbf{R2}$). For instance, in 2013 Levesque et al. [155] found that the number of illegal and questionable websites visited by a user is less related to the risk of malware infection than the number of sport or computer sites. Similarly, Bossler et al. [29] found that the time spent performing illegitimate computer activities was NOT a good predictor of malware infections. Strangely, the authors found that even higher computer skills and the adoption of careful password management failed to reduce this risk.

Many independent studies [27], [28], [155] found instead evidence that the *volume* of performed actions (e.g., the number of software installed or the number of websites visited, independently from their category) was always correlated to a higher risk. If confirmed, this finding seems to suggest that there is a systematic

risk of performing common actions – such as browsing the web or installing software – and the final risk would mainly depend on how many times these simple tasks are repeated by an individual or an organization. In other words, a possible direction is to try to model the risk of a compromise by using a *frequency-based* approach (R3), which is already a common solution to describe safety risks.

Measure the attack surface. In a given cyber environment, the attack surface is defined as the set of different points where an attacker can try to break into the system or exfiltrate information. As a direct consequence, reducing the attack surface by removing unnecessary services or limiting the access to parts of the infrastructure represents a way to increase the security by reducing the number of components that an attacker can target. The rationale behind this concept is that the likelihood of suffering from a security issue will raise according to the number and diversity of software, services, and systems used. While this is simple mathematics (and approached have been proposed to measure the attack surface of a system [156], [157]), the exact relationship that these variables have with cyber risk is still unknown and more experiments are needed to measure how risk actually reduces with the reduction of the attack surface (R4).

Influence of business sector, reputation, and assets of an organization. As we already mentioned above, non-technical characteristics of the target can influence the number, type, and sophistication of the adversaries it needs to face. Today it is widely accepted the hypothesis that, given enough time and resources, motivated attackers can always find a way to compromise a target. Large state-sponsored cyber attacks have shown this to be the case also for the most secure government organizations [158], [159]. Therefore, the type of business, the sector, the reputation, and the assets owned by an organization may influence the risk of compromise more than other technical indicators, as they allow to capture the characteristics of the *attackers* (incentives, risks, and resources as proposed by [137]) instead of those of the defender. This assumption has already been shown to be valid to characterize both the number and the type of attacks, respectively by Sarabi et al. [21] and Thonnard et al. [22]. Moreover, this approach could also cover the risk of targeted attacks, whose ad-hoc natures does not allow them to be easily described by a frequency-based model [160].

Predict future events based on historical data. Historical data about claims and incidents are routinely used to estimate the risk in other insurance sectors. However, as already stated in section III, the use of previously collected data to predict future cyber events faces several challenges. First of all, data on cyber incidents are scant and often biased towards those events whose disclosure is mandatory because regulated by law [56], [57]. A second challenge in this approach is to shed light on the so-called repeat players. Although previous studies found a systematic difference between costs incurred by companies that experience single or multiple incidents [150] (the so-called *repeat players*), it is still not clear whether having already been compromised is a good indicator of being again compromised in the future (R5). Finally, an additional complication is represented by the fact that attack techniques evolve very rapidly over time, making obsolete results obtained from the observation of old data. For instance, if a known vulnerability

associated with a high-risk factor were to be patched, past records about events occurred because of its presence would probably not provide any contribution to capture the risk associated to new attacks.

Measure the risk that propagates through third-party relations. Outsourcing many critical business operations became a norm in the last decade. It is very typical to store and process data owned by companies on third-party cloud services and even common services such as DNS and emails are now outsourced to the cloud. This largely complicates the picture for cyber insurances as it is harder to draw a clear line of the boundaries of a company. As common sense suggests, a company that is in relation to other risky entities should have higher risk itself. While constructing sufficiently accurate service-dependency graphs of businesses is a challenging research topic by itself [161], measuring the amount of risk that propagates through this graph is an open research problem that needs attention from the community. We will come back to discuss this problem in more details in Section VII.

User's weaknesses and social engineering. One of the most common techniques used today to gain access to a network or system is social engineering: indeed, while one can think that the most successful breaches are the result of technical flaws or zero-day vulnerabilities exploitations, almost 97% of them is achieved by tricking users to reveal sensitive information using a social engineering scheme [162]. Unfortunately, while social engineering attacks can pose a tremendous threat to organizations, current approaches to IT security and risk management tend to underestimate or completely ignore the human factor in risk assessment models, tools and processes [163]. Extending existing schemes by modeling users and their behavior could largely increase their prediction accuracy (R6).

Risk aggregation. All the factors we previously mentioned are likely to somehow affect (to a different and still unknown extent) the risk of cyber incidents. But even if researchers would be able to precisely identify a number of good and stable risk indicators, we would still have known very little about the aggregation procedure required to combine the different scores. This problem is exacerbated by the fact that, for practical reasons, each study looks at a single factor in isolation. But different factors are probably not independent and they can have very complex consequences and side-effects on other indicators. For instance, a good security posture may mitigate a larger attack surface, but it can be completely undermined by untrained users. Therefore, if distinct studies respectively find good predictors of risk, a constructive combination of them would still require a considerable amount of research (R7). A classic insurance solution could be to evaluate all risk indicators separately and then rely on actuarial data about past incidents to combine them in a single risk class, but as we already said this data may be very hard to put together and may become obsolete very fast. Finally, a major obstacle to risk aggregation is the different granularity of the risk computed by different approaches. Some can predict the risk of compromise of a given software artifact, other of a user, or of an individual machine. How to aggregate these values, for example, at a company level is still an open research problem (R8).

A. Horizontal Issues

So far, we discussed different open problems and research questions and their relevance for cyber insurance. However, we believe it is important to also highlight three important aspects about cyber risk itself that apply to all previously mentioned approaches:

- 1) *Cyber risk vs cyber-insurance risk*: as briefly shown in section III-B, almost all the existing literature focuses on cyber risk assessment or prediction. Although these are important for the purpose of diverting security spendings towards most relevant threats, such evaluation could be misleading for cyber-insurance risk assessment. Indeed, a quantification of the first does not necessarily reflect the second, that after all is the actual value insurances are interested in: for instance, a class of events could have a high risk to harm one entity but lead to claim submissions with a very low probability. In other words, it is also important to study and measure how cyber risks translate to insurance claims in the real world (R9).
- 2) *Consumers vs corporations*: since cyber-insurance products are recently made available also for the consumers market [164], it is possible that a different approach and/or set of features should be considered depending on the entity under investigation. Indeed, consumers are less active with respect to big corporations, operate in a different scenario, and may become an appealing target of cyber attacks for different reasons compared to large enterprises. However, no study exists to date to compare the risk and threats encountered by consumer vs enterprise users (R10).
- 3) *Risk variety*: risk assessment or prediction procedures need to be targeted towards specific categories of risk. Indeed in an insurance context, addressing cyber risk as a single-unit problem may be too generic and may not lead to meaningful results. For instance, the authors of [25], [26] predict machines and users at risk of malware infections, without providing any fine-grained categorization (after all, malware is a very generic term). In the same way, Liu et al. [20] attempt to forecast generic cyber incidents specifying no type or effect. However, as shown by Eling et al. by using actuarial data [128], different types of data breaches need to be modeled as distinct risk categories. A more fine-grained classification is needed (R11) to also highlight particular categories of threats strongly coupled to the subject we are evaluating the risk for: for instance, malware targeted against banking systems are probably not very relevant for those enterprises in other business sectors.

VI. AREA 2: AUTOMATED DATA COLLECTION

“If you’re writing policies for personal automobile or personal homeowners insurance you definitely have a lot of really good data. The worst data is probably in cyber insurance”

– Nick Economidis, Cyber liability underwriter at Beazley PLC [47]

The importance of data collection for cyber-insurance carriers does not only relate to the actuarial domain, whose issues have already been discussed in section II. Data collection about prospective

clients is indeed the first crucial task of policy underwriting, as it allows insurance firms to elicit a reasonable approximation of the overall security posture of the applicants, measure their level of risk, and subsequently compute premiums. The most common way to achieve this goal is to furnish organizations wishing to buy a cyber-insurance policy with security questionnaires. In a recent study, Romanosky et al. [37] analyze 44 of these questionnaires filed across the states of California, Pennsylvania, and New York, and point out commonalities that allow to group the questions into four macro categories.

The first set of questions aims at defining some general *organizational* details of the company, like its business sector and annual revenues, the kind of sensitive information stored and handled, how relationships with third-party service providers are managed, the nature and amount of IT security investments and, if any, its cyber-incident history. The second category focuses on *technical* aspects, often covering questions on security and access control measures adopted by the company and, less frequently, on its information technology and computing infrastructure. The existence of *policies and procedures* for data management is investigated in a third set of questions, in which insurance firms investigate whether data processing, retention and destruction practices are compliant with current regulation laws and procedures to maintain and strengthen information security. Finally on the *legal* side, questionnaires verify how well a variety of laws and regulations, enacted to protect consumers from the consequences of cyber incidents and data breaches, are implemented and adhered.

The information collected is then used for premium computation: while some carriers use flat-rate pricing for each first- and third-party coverage (with no differentiation by firm or industry), others incorporate more features (such as firm’s sector and revenue) as factors to be multiplied in a base rate pricing. In more sophisticated policies, also the soundness and completeness of security controls and practices have a weight in the final result.

Although these questionnaires are widely adopted by cyber-insurance firms, the measure of their accuracy as a standalone tool for defining the security posture —and as a step further the risk — of an organization is still questionable. A recent work examines 24 application forms to determine whether the collection of security checks referred by technical questions corresponds to the controls defined in two well-known standards of security best practices [165]. As result, existing forms are found to be predominantly focused on a small range of controls and the authors suggest how to extend them to be in alignment with the two information security frameworks. Nevertheless, the extent to which security standards compliance reflect the level of risk a company faces has not been yet understood (R12).

As suggested by modern approaches for data collection about cyber-insurance applicants [50], cyber questionnaires should be only one of many tools employed by insurance firms. For instance, instead of relying on self-assessment, the security posture of an organization can be automatically refined using two types of data sources: (i) *internal data*, provided by monitoring and telemetry tools installed inside the subject under investigation; and (ii) *external data*, collected from publicly available databases or by scanning Internet-facing services.

Although recent works show the feasibility of both ap-

proaches [20], [25], open questions still exist on both sides. Intuitively, internal data (if available) should provide a better accuracy to understand the cybersecurity risks of an organization. However, organizations do not exist in the void and the outcomes of internal telemetry analysis could be insufficient when assessing the security posture of an entity that maintains relationships or dependencies with external subjects – thus requiring a combination of the two approaches to cover unavailable information about these third parties.

On the other hand, in a cyber-insurance scenario, internal data could be unavailable to the insurer, who needs to base his evaluation on external data only. In this respect, the effectiveness of methods based on such sources only is not known, neither conditions and circumstances in which they can be used to achieve a good accuracy. As already depicted in Table II, studies that use external indicators to predict risk also validate their findings based on externally available ground truth. This is a big limitation, as cyber incidents are insufficiently reported and records, even if available, are often published too late and miss details and key elements. Moreover, the different precision and granularity of the ground truth make impossible to compare the results with those obtained with internal indicators. More research is therefore needed to compare the accuracy and relationship of external indicators and internal telemetry information on the same dataset (R13). In particular, no previous work has provided insights on a combined use of both sources, trying to answer the question whether internal data can serve as ground truth for refining the power of external indicators (R14).

VII. AREA 3: CATASTROPHE MODELING

“One key challenge is accumulation. [...] We know we can write earthquake exposures in both Japan and California with the confidence that the same event will not impact all these exposures at once. We know to be wary of writing two industrial risks along the same river basin, and the role flood defenses play in mitigating loss. With cyber risks, the contours of systemic accumulation are not as clear”

– Hemant Shah, Risk Management Solutions [59].

For an insurance company, catastrophe modeling (or simply cat modeling) is a way to estimate the likelihood or frequency at which catastrophes can occur and to what extent they can impact the insurance. To decrease the likelihood of cyber catastrophes, a typical solution that is widely adopted is client diversification. The assumption here is that if the clients of the insurance company have diverse attack surfaces and diverse characteristics, a potential new zero-day vulnerability will not exist in all of them, leaving only a percentage of insureds affected by a possible cyber attack. While this may seem a reasonable conclusion, a recent unpublished work from Eling and Schnell [166] suggested that, when modeling losses with specific distributions, diversification may not be a good idea because of the heavy-tailed distribution nature of cyber risks. This would be an important and counter-intuitive finding, that needs to be confirmed by further measurements (R15).

At its core, cat modeling boils down to capturing and modeling dependencies among different entities. This, in turn, translates into the identification of the dependencies that come from the software, hardware, and services used by a company. However, obtaining such

detailed and comprehensive information about a large enterprise is a very challenging task. Moreover, because of the cyber-insurance context and the transitive nature of these dependencies, this task would need to be performed by using publicly available datasets. This makes the problem even more complex and hence, we believe, it opens new directions for researchers to explore and contribute.

In an ideal scenario where all companies reveal the software, hardware, and services they use and share with the community, building the service dependency graph, identifying the nodes in this graph that might cause catastrophic events, and calculating the indirect risk that comes from these dependencies would be a simple task. However, even in such a perfect world, the dynamicity of the graph would require to continuously report and recalculate the risk and likelihood of the existing catastrophes and the identification of new catastrophe scenarios. In other domains, if two risks are not connected (such as a fire hazard on two areas tens of thousands of miles apart) this fact is not likely to change in the near future. But in the cyber-insurance domain, the relationships among two different companies are often very ephemeral – as services providers and software libraries may change very often. But as of now, there is no existing work that studied how the dynamicity of the ecosystem could influence the whole cat modeling topic and whether (and how often) the portfolios defined by the insurance companies should also be updated (R16).

Moreover, the reality is far from this ideal scenario and even the topic of building adequately accurate service dependency graphs and modeling the catastrophes with sparse and incomplete data are research topics that need more attention from the community (R17). Altogether, this can lead to a *supply chain risk analysis* that would provide a principled foundation for catastrophe modeling.

However, the identification of all services used by a company, especially without its cooperation, is often infeasible. For instance, the presence of backup or redundancy services can remain undetected, as those only come into play when the primary provider fails. As a full and precise view of all the dependencies of a company may be impossible to obtain, then a modeling algorithm should be able to work with incomplete information, potentially inferring the missing connections from settings and relations observed elsewhere (R18). Although not done particularly for the cyber-insurance domain, there exist two works [161], [167] that aimed at building dependency graphs of popular companies by using public datasets such as RIPE atlas, passive DNS records, and web crawling data. In 2017, Dell’amico et al [161] performed a large-scale study to identify the dependencies between websites and Internet services. The findings of the study confirm the monopoly problem in the current service ecosystem. To make matters worse, over time the Internet appears to be losing its decentralized nature and the popularity of the few dominant providers is steadily increasing. In the same year, Simeonovski et al [167] built a service dependency graph to explore what percentage of the Internet would be affected when a popular provider is attacked. The study found that by only targeting a handful of service providers it would be possible to take down 23% of the websites.

Another challenge that affects cat modeling is the lack of a mapping procedure to reliably associate measurements and public data to organizations. Network scans, web crawlers, service monitoring systems, public blacklist, and other techniques that can be used to identify the software and technologies adopted

by a company typically work at the level of domain names or IP addresses. On the contrary, incident reports and risk prediction operate at a company granularity. Sadly, the connection between the two is not always straightforward and new techniques are needed to link the two information (R19). For instance, Liu et al. [20] explain their attempt to perform a manual mapping and all the difficulties and caveats encountered in the process, making it evident the necessity of a clearer and automated procedure.

VIII. AREA 4: FORENSIC ANALYSIS

“I often think of the 1990s as the decade of prevention, the 2000 as the decade of detection, and this is the decade of incident response.”

– Bruce Schneier, Security Specialist

After the detection of a cyber incident, the response phase requires the intervention of computer security experts to analyze and understand the detail of the event. However, computer security skills are not only required for helping the company to recover from the incident but also, from an insurer’s perspective, to verify the claim, assess the damage, and confirm whether it is covered by the subscriber’s policy. Indeed, forensic investigations are the norm to assess if, and to which extent, the insurance is liable for the event.

Computer forensics is a broad research field that covers the collection, analysis, and preservation of digital evidence. It is a highly developed science with its own language, modus operandi, and standardized procedures [168]. However, while the other research topics discussed in this paper have all been recently contextualized (in terms of specific problems and new challenges) to the cyber-insurance domain, no study has looked at the problem of computer forensics from a cyber-insurance perspective.

For instance, one aspect that may require special attention is information forgery. In traditional insurance sectors, fake accidents cost over 30 billion dollars per year, with several insurers reporting these frauds to account for up to 20% of claims costs [169]. However, while set-up wrecks and burning houses are sadly common practice for fraudsters to cash the insurance coverages, there is almost no mention to date about similar frauds in the cyber domain.

Current forensic approaches are mainly concerned with the possibility that an attacker can hide undetected or that important evidence and artifacts can be deleted or manipulated. In other words, the focus on evasion and not on forgery. The lack of motivation can explain why planting fake evidence in a computer system is not yet very common, but forged incidents are extremely easy to set up for anyone with average programming skills [170]. The vast majority of the indicator of compromise used today rely on the simple existence of filesystem and registry artifacts - without any knowledge of how (and by whom) the data was created in the first place. In this setting, it is not hard to mimic a malware infection or even a targeted attack against an organization. However, with cyber insurances becoming more and more common, forged digital evidence may become a major problem in the future.

In particular, digital evidence forgery could help businesses to overcome one of the cyber-insurance most-common pitfalls: the fact that technicalities can invalidate coverage allowing insurance carriers to deny indemnity payments [171]. For instance, cyber insurance does not normally cover when employee errors (e.g.,

falling for phishing attacks) are the cause of a malware infection (e.g., ransomware) [172]. Since these events are instead covered under other clauses (e.g., malware installed by an external attacker), forging digital evidence would allow to “fake” a botnet infection to fall within the scenario covered by the insurance policy, thus allowing the victim to cash the indemnity.

Since today staging fake security incidents requires very little effort, researchers should not only study how to collect hidden signs of compromise, but also how to double-check and validate their authenticity (R20).

IX. CONCLUSIONS

In this paper we discussed the unique challenges that affect the cyber-insurance sector. We focus on a pure technical perspective, highlighting the limitations of current approaches, evaluating the feasibility of new solutions, and proposing research areas in which system and network security experts can play a fundamental role for the development of cyber insurance. Differently from legacy frameworks based on qualitative approaches for risk assessment and data collection, we endorse the relevance of prediction techniques based on objective measures and automatic feature gathering.

ACKNOWLEDGMENT

This project was supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 771844 BitCrumbs) and by the European Unions Horizon 2020 research and innovation programme under grant agreement No. 786669 (ReAct).

REFERENCES

- [1] “Cyber Incident & Breach Trends Report,” tech. rep., Internet Society (ISOC), Jan. 2018. Available at: https://www.otalliance.org/system/files/files/initiative/documents/ota_cyber_incident_trends_report_jan2018.pdf.
- [2] BDO, “Bdo - board survey,” tech. rep., BDO - USA, 2016. Available at: <https://www.bdo.com/insights/assurance/client-advisories/2016-board-survey>.
- [3] “Managing Cyber Security as a Business Risk: Cyber Insurance in the Digital Age,” tech. rep., Ponemon institute, Aug. 2013. Available at: <https://www.ponemon.org/local/upload/file/Cyber%20Insurance%20white%20paper%20FINAL%207.pdf>.
- [4] Gartner, “Gartner forecasts worldwide information security spending to exceed \$124 billion in 2019,” <https://gtr.it/20lkhYz>. Accessed: 2019-07-31.
- [5] R. Böhme, G. Schwartz, et al., “Modeling cyber-insurance: Towards a unifying framework,” in *WEIS*, 2010.
- [6] “Global cyber market overview,” tech. rep., Aon Inpoint, June 2017. Available at: <https://www.aon.com/inpoint/bin/pdfs/white-papers/Cyber.pdf>.
- [7] “Cyber insurance claims: Ransomware disrupts business,” tech. rep., American International Group AIG, Mar. 2018. Available at: <https://www.aig.co.uk/content/dam/aig/emea/united-kingdom/documents/Insights/cyber-claims-report-may-18.pdf>.
- [8] “insurers pay large claims for high profile cyber attacks,” tech. rep., Jardine Lloyd Thompson group plc, June 2018. Available at: <https://www.jlt.com/insurance-risk/cyber-insurance/insights/insurers-pay-large-claims-for-high-profile-cyber-attacks>.
- [9] M. Potdar, “Cyber insurance market overview,” <https://www.alliedmarketresearch.com/cyber-insurance-market>. Accessed: 2019-07-31.
- [10] P. Nohe, “Re-hashed: 2018 cybercrime statistics: A closer look at the web of profit,” <https://www.thesslstore.com/blog/2018-cybercrime-statistics/>. Accessed: 2019-07-31.
- [11] B. Sussman, “5 reasons cyber insurance market will hit \$23 billion,” <https://www.secureworldexpo.com/industry-news/5-reasons-cyber-insurance-market-will-hit-23-billion>. Accessed: 2019-07-31.
- [12] P. Nohe, “Study: 66% of smbs would shut down if hit by a data breach,” <https://www.thesslstore.com/blog/study-66-smbs-shut-hit-data-breach>. Accessed: 2019-07-31.

- [13] S. White, "What is GDPR's impact on Cyber Insurance?," <https://gdpr-report/news/2019/05/15/what-is-the-impact-of-gdprs-on-cyber-insurance>. Accessed: 2019-07-31.
- [14] "Cyber Risk: Too Big to Insure?," tech. rep., Institute of Insurance Economics and Swiss Re, June 2016. Available at: <https://www.ivw.unisg.ch/~media/internet/content/dateien/instituteundcenters/ivw/studien/cyberriisk2016.pdf>.
- [15] "Insurability of cyber risk," tech. rep., The Geneva Associations, Aug. 2014. Available at: https://www.genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public/ga2014-if14-biener_elingwirfs.pdf.
- [16] C. Biener, M. Eling, and J. H. Wirfs, "Insurability of cyber risk: An empirical analysis," *The Geneva Papers on Risk and Insurance-Issues and Practice*, vol. 40, no. 1, pp. 131–158, 2015.
- [17] "Views from the C-suite Survey 2018," tech. rep., FICO, Mar. 2018. Available at: <https://www.fico.com/en/resource-download-file/6341>.
- [18] R. P. Majuca, W. Yurcik, and J. P. Kesan, "The evolution of cyberinsurance," *arXiv preprint cs/0601020*, 2006.
- [19] A. Marotta, F. Martinelli, S. Nanni, A. Orlando, and A. Yautsiukhin, "Cyber-insurance survey," *Computer Science Review*, vol. 24, pp. 35–61, 2017.
- [20] Y. Liu, A. Sarabi, J. Zhang, P. Naghizadeh, M. Karir, M. Bailey, and M. Liu, "Cloudy with a chance of breach: Forecasting cyber security incidents," in *USENIX Security Symposium*, pp. 1009–1024, 2015.
- [21] A. Sarabi, P. Naghizadeh, Y. Liu, and M. Liu, "Prioritizing security spending: A quantitative analysis of risk distributions for different business profiles,"
- [22] O. Thonnard, L. Bilge, A. Kashyap, and M. Lee, "Are you at risk? profiling organizations and individuals subject to targeted attacks," in *International Conference on Financial Cryptography and Data Security*, pp. 13–31, Springer, 2015.
- [23] H. K. Browne, W. A. Arbaugh, J. McHugh, and W. L. Fithen, "A trend analysis of exploitations," in *Security and Privacy, 2001. S&P 2001. Proceedings. 2001 IEEE Symposium on*, pp. 214–229, IEEE, 2001.
- [24] M. Bozorgi, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond heuristics: learning to classify vulnerabilities and predict exploits," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 105–114, ACM, 2010.
- [25] L. Bilge, Y. Han, and M. Dell'Amico, "Riskteller: Predicting the risk of cyber incidents," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1299–1311, ACM, 2017.
- [26] T.-F. Yen, V. Heorhiadi, A. Oprea, M. K. Reiter, and A. Juels, "An epidemiological study of malware encounters in a large enterprise," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1117–1130, ACM, 2014.
- [27] D. Canali, L. Bilge, and D. Balzarotti, "On the effectiveness of risk prediction based on users browsing behavior," in *Proceedings of the 9th ACM symposium on Information, computer and communications security*, pp. 171–182, ACM, 2014.
- [28] M. Sharif, J. Urakawa, N. Christin, A. Kubota, and A. Yamada, "Predicting impending exposure to malicious content from user behavior," 2018.
- [29] A. M. Bossler and T. J. Holt, "On-line activities, guardianship, and malware infection: An examination of routine activities theory," *International Journal of Cyber Criminology*, vol. 3, no. 1, 2009.
- [30] L. Allodi and F. Massacci, "Security events and vulnerability data for cyber-security risk estimation," *Risk Analysis*, vol. 37, no. 8, pp. 1606–1627, 2017.
- [31] E. J. Vaughan and T. Vaughan, *Fundamentals of risk and insurance*. John Wiley & Sons, 2007.
- [32] G. Stoneburner, A. Y. Goguen, and A. Feringa, "Sp 800-30. risk management guide for information technology systems," 2002.
- [33] M. Amutio, J. Candau, and J. Mañas, "Magerit-version 3, methodology for information systems risk analysis and management, book i-the method," *Ministerio de administraciones públicas*, 2014.
- [34] C. J. Alberts and A. Dorofee, *Managing information security risks: the OCTAVE approach*. Addison-Wesley Longman Publishing Co., Inc., 2002.
- [35] B. Karabacak and I. Sogukpinar, "Isram: information security risk analysis method," *Computers & Security*, vol. 24, no. 2, pp. 147–159, 2005.
- [36] M. CLUSIF, "Processing guide for risk analysis and management," *Club De La Securite De L'information Francias, 2nd Edition (April 2011)*, 2010.
- [37] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones, "Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk?," 2017.
- [38] U. Franke, "The cyber insurance market in sweden," *Computers & Security*, vol. 68, pp. 130–144, 2017.
- [39] "The iso/iec 27000 family of information security standards," <https://www.itgovernance.co.uk/iso27000-family#4>. Accessed: 2019-07-31.
- [40] T. Bandyopadhyay, V. S. Mookerjee, and R. C. Rao, "Why it managers don't go for cyber-insurance products," *Communications of the ACM*, vol. 52, no. 11, pp. 68–73, 2009.
- [41] N. Robinson, "Incentives and barriers of the cyber insurance market in europe," 2012.
- [42] N. Shetty, G. Schwartz, M. Felegyhazi, and J. Walrand, "Competitive cyber-insurance and internet security," in *Economics of information security and privacy*, pp. 229–247, Springer, 2010.
- [43] L. A. Gordon, M. Loeb, and T. Sohal, "A framework for using insurance for cyber-risk management," vol. 46, pp. 81–85, 03 2003.
- [44] A. Hedrick, "Cyberinsurance: a risk management tool?," in *Proceedings of the 4th annual conference on Information security curriculum development*, p. 20, ACM, 2007.
- [45] C. Toregas and N. Zahn, "Insurance for cyber attacks: The issue of setting premiums in context," *George Washington University*, 2014.
- [46] P. H. Meland, I. A. Tøndel, M. Moe, and F. Seehusen, "Facing uncertainty in cyber insurance policies," in *International Workshop on Security and Trust Management*, pp. 89–100, Springer, 2017.
- [47] J. Wolff, "Cyberinsurance tackles the wildly unpredictable world of hacks," <https://www.wired.com/story/cyberinsurance-tackles-the-wildly-unpredictable-world-of-hacks/>. Accessed: 2019-07-31.
- [48] R. Böhme and G. Kataria, "On the limits of cyber-insurance," in *International Conference on Trust, Privacy and Security in Digital Business*, pp. 31–40, Springer, 2006.
- [49] Wikipedia, "2016 dyn cyberattack: Timeline and impact," https://en.wikipedia.org/wiki/2016_Dyn_cyberattack. Accessed: 2019-07-31.
- [50] A. Caglar, "A new approach to risk assessment for cyber insurance," <https://technet.microsoft.com/en-us/library/cc163143.aspx>. Accessed: 2019-07-31.
- [51] L. Heslault, "Actuaries beware: Pricing cyber insurance is a different ballgame," <https://www.linkedin.com/pulse/actuaries-beware-pricing-cyber-insurance-different-laurent-heslault>. Accessed: 2019-07-31.
- [52] Gemalto, "Cyber insurance: The challenges facing actuaries in measuring cyber risk," <https://blog.gemalto.com/security/2017/10/17/cyber-insurance-challenges-facing-actuaries-measuring-cyber-risk/>. Accessed: 2019-07-31.
- [53] "Hackmageddon: Information security timelines and statistics," <https://www.hackmageddon.com>. Accessed: 2019-07-31.
- [54] "The veris community database (vcdb)," <http://veriscommunity.net/vcdb.html>. Accessed: 2019-07-31.
- [55] "Owasp wasc web hacking incidents database project," https://www.owasp.org/index.php/OWASP_WASC_Web_Hacking_Incidents_Database_Project. Accessed: 2019-07-31.
- [56] "Security breach notification laws," <http://www.ncsl.org/research/telecommunications-and-information-technology/security-breach-notification-laws.aspx>. Accessed: 2019-07-31.
- [57] "An overview of the main changes under gdpr and how they differ from the previous directive," <https://www.eugdpr.org/key-changes.html>. Accessed: 2019-07-31.
- [58] A. R. Tomas Girmius, Scott Stransky, "Aggregated cyber risk: The nightmare scenarios," <http://www.air-worldwide.com/Publications/AIR-Currents/2015/Aggregated-Cyber-Risk--The-Nightmare-Scenarios/>. Accessed: 2019-07-31.
- [59] I. Cambridge Centre for Risk Studies & Risk Management Solutions, "Managing cyber insurance accumulation risk," https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-rms-managing-cyber-insurance-accumulation-risk.pdf. Accessed: 2019-07-31.
- [60] I. Cambridge Centre for Risk Studies & Risk Management Solutions, "Sybil logic bomb cyber catastrophe scenario," https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-sybil-logic-bomb-cyber-catastrophe-stress-test.pdf. Accessed: 2019-07-31.
- [61] D. Geer, R. Bace, P. Gutmann, P. Metzger, C. P. Pfleeger, J. S. Quarterman, and B. Schneier, "Cyberinsecurity: The cost of monopoly," *Computer and Communications Industry Association (CCIA)*, 2003.
- [62] K. P. Birman and F. B. Schneider, "The monoculture risk put into context," *IEEE Security & Privacy*, vol. 7, no. 1, pp. 14–17, 2009.
- [63] M. Stamp, "Risks of monoculture," *Communications of the ACM*, vol. 47, no. 3, p. 120, 2004.
- [64] B. Schneier, "The dangers of a software monoculture," *Information Security Magazine*, 2010.
- [65] K. J. Higgins, "More than a half million servers exposed to heartbleed flaw," <https://www.darkreading.com/informationweek-home/more-than-a-half-million-servers-exposed-to-heartbleed-flaw/d/d-id/1204318>. Accessed: 2019-07-31.

- [66] C. Wright, "Understanding kaminsky's dns bug." <http://www.linuxjournal.com/content/understanding-kaminskys-dns-bug>. Accessed: 2019-07-31.
- [67] P. Nickinson, "The 'stagefright' exploit: What you need to know." <https://www.androidcentral.com/stagefright>. Accessed: 2019-07-31.
- [68] S. Corporation, "Ssl 3.0 poodle attack vulnerability." https://support.symantec.com/en_US/article.TECH226102.html. Accessed: 2019-07-31.
- [69] M. S. Bulletin, "Vulnerability in server service could allow remote code execution." <https://docs.microsoft.com/en-us/security-updates/SecurityBulletins/2008/ms08-067>. Accessed: 2019-07-31.
- [70] J. Williams, "Java serialization vulnerability threatens millions of applications." <https://www.contrastsecurity.com/security-influencers/java-serialization-vulnerability-threatens-millions-of-applications>. Accessed: 2019-07-31.
- [71] P. Kocher, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," *arXiv preprint arXiv:1801.01203*, 2018.
- [72] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown," *arXiv preprint arXiv:1801.01207*, 2018.
- [73] L. Finance, "Promoting uk cyber prosperity: Public-private cyber-catastrophe reinsurance," 2015.
- [74] Artemis, "Market unprepared for silent cyber loss aggregation: Welsh, sciemus." <http://www.artemis.bm/blog/2016/12/01/market-unprepared-for-silent-cyber-loss-aggregation-welsh-sciemus/>.
- [75] R. news, "Cyber re/insurance market 'frustratingly immature'." <https://www.reinsurancencs.ws/cyber-re-insurance-market-frustratingly-immature-inga-beale-lloyds/>.
- [76] A. B. Association and FSSCC, "Cyber insurance buying guide." https://www.aba.com/Tools/Function/Documents/2016Cyber-Insurance-Buying-Guide_FINAL.pdf. Accessed: 2019-07-31.
- [77] L. Bailey, "Mitigating moral hazard in cyber-risk insurance," *JL & Cyber Warfare*, vol. 3, p. 1, 2014.
- [78] I. A. Tøndel, P. H. Meland, A. Omerovic, E. A. Gjære, and B. Solhaug, "Using cyber-insurance as a risk management strategy: Knowledge gaps and recommendations for further research," 2015.
- [79] J. P. Kesan, R. P. Majuca, and W. J. Yurcik, "The economic case for cyberinsurance," 2004.
- [80] H. ulisi Ogut and S. Raghunathan, "Cyber insurance and it security investment: Impact of interdependent risk,"
- [81] X. Zhao, L. Xue, and A. B. Whinston, "Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling," *ICIS 2009 Proceedings*, p. 49, 2009.
- [82] W. Shim, "An analysis of information security management strategies in the presence of interdependent security risk," *Asia Pacific Journal of Information Systems*, vol. 22, no. 1, pp. 79–101, 2012.
- [83] M. Lelarge and J. Bolot, "Economic incentives to increase security in the internet: The case for insurance," in *INFOCOM 2009, IEEE*, pp. 1494–1502, IEEE, 2009.
- [84] J.-C. Bolot and M. Lelarge, "A new perspective on internet security using insurance," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pp. 1948–1956, IEEE, 2008.
- [85] F. Martinelli, A. Orlando, G. Uganbayar, and A. Yautsiukhin, "Preventing the drop in security investments for non-competitive cyber-insurance market," in *International Conference on Risks and Security of Internet and Systems*, pp. 159–174, Springer, 2017.
- [86] J. Bolot and M. Lelarge, "Cyber insurance as an incentive for internet security," in *Managing information risk and the economics of security*, pp. 269–290, Springer, 2009.
- [87] N. Shetty, G. Schwartz, and J. Walrand, "Can competitive insurers improve network security?," in *International Conference on Trust and Trustworthy Computing*, pp. 308–322, Springer, 2010.
- [88] G. A. Schwartz and S. S. Sastry, "Cyber-insurance framework for large scale interdependent networks," in *Proceedings of the 3rd international conference on High confidence networked systems*, pp. 145–154, ACM, 2014.
- [89] R. Pal, L. Golubchik, K. Psounis, and P. Hui, "Will cyber-insurance improve network security? a market analysis," in *INFOCOM, 2014 Proceedings IEEE*, pp. 235–243, IEEE, 2014.
- [90] A. Laszka and J. Grossklags, "Should cyber-insurance providers invest in software security?," in *European Symposium on Research in Computer Security*, pp. 483–502, Springer, 2015.
- [91] G. Schwartz, N. Shetty, and J. Walrand, "Cyber-insurance: Missing market driven by user heterogeneity," *preparation, www.eecs.berkeley.edu/nikhils/SecTypes.pdf*, 2010.
- [92] R. Pal and L. Golubchik, "On the economics of information security: the problem of designing optimal cyber-insurance contracts," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 2, pp. 51–53, 2010.
- [93] Z. Yang and J. C. Lui, "Security adoption and influence of cyber-insurance markets in heterogeneous networks," *Performance Evaluation*, vol. 74, pp. 1–17, 2014.
- [94] P. Naghizadeh and M. Liu, "Voluntary participation in cyber-insurance markets," in *Workshop on the Economics of Information Security (WEIS)*, 2014.
- [95] "Information technology - Security techniques - Information security risk management," standard, International Organization for Standardization, June 2011.
- [96] "Isaca - cobit." <http://www.isaca.org/COBIT/Pages/default.aspx>. Accessed: 2019-07-31.
- [97] "Microsoft solutions for security and compliance and microsoft security center of excellence - the security risk management guide." <https://technet.microsoft.com/en-us/library/cc163143.aspx>. Accessed: 2019-07-31.
- [98] C. club da la sécurité de l'information français, "Risk management: concept and methods," tech. rep., 2009.
- [99] R. A. Caralli, J. F. Stevens, L. R. Young, and W. R. Wilson, "Introducing octave allegro: Improving the information security risk assessment process," tech. rep., CARNEGIE-MELLON UNIV PITTSBURGH PA SOFTWARE ENGINEERING INST, 2007.
- [100] I. I. E. Commission, "Iec bs en 61882:2016 - hazard and operability study (hazop studies) - application guide," tech. rep., 2016.
- [101] "Guide for conducting risk assessments," tech. rep., NIST - National Institute of Standards and Technology, 2012.
- [102] M. S. Lund, B. Solhaug, and K. Stølen, *Model-driven risk analysis: the CORAS approach*. Springer Science & Business Media, 2010.
- [103] S. A. Butler, "Security attribute evaluation method: a cost-benefit approach," in *Proceedings of the 24th international conference on Software engineering*, pp. 232–240, ACM, 2002.
- [104] J. H. Pardue and P. Patidar, "Thrats to healthcare date: A threat tree for risk assessment," *Issues in Information Systems*, vol. 12, no. 1, pp. 106–113, 2011.
- [105] E. J. Byres, M. Franz, and D. Miller, "The use of attack trees in assessing vulnerabilities in scada systems," in *Proceedings of the international infrastructure survivability workshop*, 2004.
- [106] B. Kordy, S. Mauw, S. Radomirović, and P. Schweitzer, "Foundations of attack-defense trees," in *International Workshop on Formal Aspects in Security and Trust*, pp. 80–95, Springer, 2010.
- [107] C.-W. Ten, C.-C. Liu, and M. Govindarasu, "Vulnerability assessment of cybersecurity for scada systems using attack trees," in *Power Engineering Society General Meeting, 2007. IEEE*, pp. 1–8, IEEE, 2007.
- [108] I. Ray and N. Poolsapassit, "Using attack trees to identify malicious attacks from authorized insiders," in *European Symposium on Research in Computer Security*, pp. 231–246, Springer, 2005.
- [109] D. Balzarotti, M. Monga, and S. Sicari, "Assessing the risk of using vulnerable components," in *Quality of Protection: Security Measurements and Metrics (QoP)* (D. Gollmann, F. Massacci, and A. Yautsiukhin, eds.), Advances in Information Security, pp. 65–78, Springer, 2006.
- [110] O. Sheyner and J. Wing, "Tools for generating and analyzing attack graphs," in *International Symposium on Formal Methods for Components and Objects*, pp. 344–371, Springer, 2003.
- [111] N. Poolsapassit, R. Dewri, and I. Ray, "Dynamic security risk management using bayesian attack graphs," *IEEE Transactions on Dependable and Secure Computing*, vol. 9, no. 1, pp. 61–74, 2012.
- [112] S. Noel, S. Jajodia, L. Wang, and A. Singhal, "Measuring security risk of networks using attack graphs," *International Journal of Next-Generation Computing*, vol. 1, no. 1, pp. 135–147, 2010.
- [113] O. Sheyner, J. Haines, S. Jha, R. Lippmann, and J. M. Wing, "Automated generation and analysis of attack graphs," in *Security and privacy. 2002. Proceedings. 2002 IEEE Symposium on*, pp. 273–284, IEEE, 2002.
- [114] I. I. E. Commission, "Analysis techniques for system reliability - procedure for failure mode and effects analysis (fmea)," tech. rep., 2006.
- [115] P. A. Ralston, J. H. Graham, and J. L. Hieb, "Cyber security risk assessment for scada and dcs networks," *ISA transactions*, vol. 46, no. 4, pp. 583–594, 2007.
- [116] A. Cook, R. Smith, L. Maglaras, and H. Janicke, "Measuring the risk of cyber attack in industrial control systems," BCS eWiC, 2016.
- [117] C. Schmittner, T. Gruber, P. Puschner, and E. Schoitsch, "Security application of failure mode and effect analysis (fmea)," in *International Conference on Computer Safety, Reliability, and Security*, pp. 310–325, Springer, 2014.
- [118] "Cybersecurity stress-testing: Don't stress about your companys safety." <https://axiomcyber.com/cybersecurity/cybersecurity-stress-testing-dont-stress-about-your-companys-safety/>.
- [119] "Anatomy of a cyber risks stress test." <https://www.zurich.com/en/knowledge/articles/2016/11/anatomy-of-a-cyber-risks-stress-test>. Accessed: 2019-07-31.
- [120] "Cyber risk stress testing." <https://www.firstbase.co.uk/cyber-resilience/cyber-risks-stress-testing/>.

- [121] "Insurers told to conduct stress tests for cyber attacks." <https://www.ft.com/content/92ec137a-6185-11e7-8814-0ac7eb84e5f1>. Accessed: 2019-07-31.
- [122] "Cybersecurity incident simulation exercises." [https://www.ey.com/Publication/vwLUAssets/EY_-_Cybersecurity_Incident_Simulation_Exercises/\\$FILE/EY-cybersecurity-incident-simulation-exercises-scored.pdf](https://www.ey.com/Publication/vwLUAssets/EY_-_Cybersecurity_Incident_Simulation_Exercises/$FILE/EY-cybersecurity-incident-simulation-exercises-scored.pdf). Accessed: 2019-07-31.
- [123] "How to stress test your cyber risk management." <https://www.marsh.com/uk/insights/risk-in-context/how-to-stress-test-your-cyber-risk-management.html>. Accessed: 2019-07-31.
- [124] "Financial catastrophe research and stress test scenarios." https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/2016risksymposium-riskculture-slides-skelton.pdf. Accessed: 2019-07-31.
- [125] I. R. C. of the International Actuarial Association *et al.*, "Stress testing and scenario analysis." http://www.actuaries.org/CTTEES_SOLV/Documents/StressTestingPaper.pdf, 2013. Accessed: 2019-07-31.
- [126] "2017 Cyber Risk Landscape," tech. rep., Risk Management Solutions, Inc and Cambridge Centre for Risk Studies, 2017. Available at: https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/risk/downloads/crs-rms-cyber-risk-landscape-2017.pdf.
- [127] H. Herath and T. Herath, "Copula-based actuarial model for pricing cyber-insurance policies," *Insurance Markets and Companies: Analyses and Actuarial Computations*, vol. 2, no. 1, pp. 7–20, 2011.
- [128] M. Eling and N. Loperfido, "Data breaches: Goodness of fit, pricing, and risk measurement," *Insurance: Mathematics and Economics*, vol. 75, pp. 126–136, 2017.
- [129] S. Wheatley, T. Maillart, and D. Sornette, "The extreme risk of personal data breaches and the erosion of privacy," *The European Physical Journal B*, vol. 89, no. 1, p. 7, 2016.
- [130] R. Layton and P. A. Watters, "A methodology for estimating the tangible cost of data breaches," *Journal of Information Security and Applications*, vol. 19, no. 6, pp. 321–330, 2014.
- [131] A. A. Yayla and Q. Hu, "The impact of information security events on the stock value of firms: The effect of contingency factors," *Journal of Information Technology*, vol. 26, no. 1, pp. 60–77, 2011.
- [132] "ISO 31000:2018 - Risk management: Principles and guidelines," standard, International Organization for Standardization, Feb. 2018.
- [133] J. Grossklags, N. Christin, and J. Chuang, "Secure or insure?: a game-theoretic analysis of information security games," in *Proceedings of the 17th international conference on World Wide Web*, pp. 209–218, ACM, 2008.
- [134] B. Johnson, R. Böhme, and J. Grossklags, "Security games with market insurance," in *International Conference on Decision and Game Theory for Security*, pp. 117–130, Springer, 2011.
- [135] H. Ogut, N. Menon, and S. Raghunathan, "Cyber insurance and it security investment: Impact of interdependence risk," in *WEIS*, 2005.
- [136] "Eiopa insurance stress tests to assess cyber risk response." <https://www.out-law.com/en/articles/2018/may/eiopa-insurance-stress-tests-cyber-risk/>.
- [137] S. E. Schechter, "Toward econometric models of the security risk from remote attack," *IEEE security & privacy*, no. 1, pp. 40–44, 2005.
- [138] S. Hakim, G. F. Rengert, and Y. Shachamurove, *Knowing your odds: Home burglary and the odds ratio*. University of Pennsylvania, Center for Analytic Research in Economics and the Social Sciences, 2000.
- [139] E. Condon, M. Cukier, and T. He, "Applying software reliability models on security incidents," in *Software Reliability, 2007. ISSRE'07. The 18th IEEE International Symposium on*, pp. 159–168, IEEE, 2007.
- [140] E. Condon, A. He, and M. Cukier, "Analysis of computer security incident data using time series models," in *Software Reliability Engineering, 2008. ISSRE 2008. 19th International Symposium on*, pp. 77–86, IEEE, 2008.
- [141] B. Edwards, S. Hofmeyr, and S. Forrest, "Hype and heavy tails: A closer look at data breaches," *Journal of Cybersecurity*, vol. 2, no. 1, pp. 3–14, 2016.
- [142] T. Maillart and D. Sornette, "Heavy-tailed distribution of cyber-risks," *The European Physical Journal B*, vol. 75, no. 3, pp. 357–364, 2010.
- [143] T. Zimmermann, N. Nagappan, and L. Williams, "Searching for a needle in a haystack: Predicting security vulnerabilities for windows vista," in *Software Testing, Verification and Validation (ICST), 2010 Third International Conference on*, pp. 421–428, IEEE, 2010.
- [144] R. Scandariato, J. Walden, A. Hovsepian, and W. Joosen, "Predicting vulnerable software components via text mining," *IEEE Transactions on Software Engineering*, vol. 40, no. 10, pp. 993–1006, 2014.
- [145] Y. Shin and L. Williams, "An empirical model to predict security vulnerabilities using code complexity metrics," in *Proceedings of the Second ACM-IEEE international symposium on Empirical software engineering and measurement*, pp. 315–317, ACM, 2008.
- [146] J. Walden, J. Stuckman, and R. Scandariato, "Predicting vulnerable components: Software metrics vs text mining," in *Software Reliability Engineering (ISSRE), 2014 IEEE 25th International Symposium on*, pp. 23–33, IEEE, 2014.
- [147] A. E. Hassan, "Predicting faults using the complexity of code changes," in *Proceedings of the 31st International Conference on Software Engineering*, pp. 78–88, IEEE Computer Society, 2009.
- [148] O. H. Alhazmi, Y. K. Malaiya, and I. Ray, "Measuring, analyzing and predicting security vulnerabilities in software systems," *Computers & Security*, vol. 26, no. 3, pp. 219–228, 2007.
- [149] J. Jacobs, "Analyzing ponemon cost of data breach," *Data Driven Security*, vol. 11, 2014.
- [150] S. Romanosky, "Examining the costs and causes of cyber incidents," *Journal of Cybersecurity*, vol. 2, no. 2, pp. 121–135, 2016.
- [151] P. Radanliev, D. De Roure, S. Cannady, R. M. Montalvo, R. Nicolescu, and M. Huth, "Economic impact of iot cyber risk-analysing past and present to predict the future developments in iot risk analysis and iot cyber insurance," 2018.
- [152] "Cyberattack insurance a challenge for business." <https://www.nytimes.com/2014/06/09/business/cyberattack-insurance-a-challenge-for-business.html>.
- [153] D. W. Straub Jr, "Effective is security: An empirical study," *Information Systems Research*, vol. 1, no. 3, pp. 255–276, 1990.
- [154] V. Verendel, "Quantified security is a weak hypothesis: a critical survey of results and assumptions," in *Proceedings of the 2009 workshop on New security paradigms workshop*, pp. 37–50, ACM, 2009.
- [155] F. Lalonde Levesque, J. Nsiempba, J. M. Fernandez, S. Chiasson, and A. Somayaji, "A clinical study of risk factors related to malware infections," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pp. 97–108, ACM, 2013.
- [156] P. K. Manadhata and J. M. Wing, "An attack surface metric," *IEEE Transactions on Software Engineering*, no. 3, pp. 371–386, 2010.
- [157] M. Howard, J. Pincus, and J. M. Wing, "Measuring relative attack surfaces," in *Computer security in the 21st century*, pp. 109–137, Springer, 2005.
- [158] "What is stuxnet, who created it and how does it work?." <https://www.csoonline.com/article/3218104/malware/what-is-stuxnet-who-created-it-and-how-does-it-work.html>.
- [159] "Inside the cunning, unprecedented hack of ukraine's power grid." <https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/>.
- [160] W. Pieters, Z. Lukszo, D. Hadziomanovic, and J. van den Berg, "Reconciling malicious and accidental risk in cyber security," *J. Internet Serv. Inf. Secur.*, vol. 4, no. 2, pp. 4–26, 2014.
- [161] M. Dell'Amico, L. Bilge, A. Kayyoor, P. Efstathiopoulos, and P.-A. Vervier, "Lean on me: Mining internet service dependencies from large-scale dns data," in *Proceedings of the 33rd Annual Computer Security Applications Conference, ACSAC 2017, (New York, NY, USA)*, pp. 449–460, ACM, 2017.
- [162] "2017 Data Breach Investigations Report," tech. rep., Verizon Enterprise, Apr. 2017. Available at: <https://www.ictsecuritymagazine.com/wp-content/uploads/2017-Data-Breach-Investigations-Report.pdf>.
- [163] R. Puricelli, "The underestimated social engineering threat in it security governance and management." <https://www.isaca.org/Journal/archives/2015/Volume-3/Pages/the-underestimated-social-engineering-threat.aspx>.
- [164] "Hsb introduces first cyber insurance for consumers." <https://www.munichre.com/HSB/first-personal-cyber-insurance/index.html>. Accessed: 2019-07-31.
- [165] D. Woods, I. Agrafiotis, J. R. Nurse, and S. Creese, "Mapping the coverage of security controls in cyber insurance proposal forms," *Journal of Internet Services and Applications*, vol. 8, no. 1, p. 8, 2017.
- [166] M. Eling and W. Schnell, "Extreme cyber risks and the non-diversification trap,"
- [167] M. Simeonovski, G. Pellegrino, C. Rossow, and M. Backes, "Who controls the internet?: Analyzing global threats using property graph traversals," in *Proceedings of the 26th International Conference on World Wide Web, WWW '17, (Republic and Canton of Geneva, Switzerland)*, pp. 647–656, International World Wide Web Conferences Steering Committee, 2017.
- [168] E. Casey, *Digital evidence and computer crime: Forensic science, computers, and the internet*. Academic press, 2011.
- [169] B. Bailey, "Fake accidents cost insurance industry billions each year, experts say." <https://newsok.com/article/5518524/fake-accidents-cost-insurance-industry-billions-each-year-experts-say>.
- [170] N. Gelernter, Y. Grinstein, and A. Herzberg, "Cross-site framing attacks," in *Proceedings of the 31st Annual Computer Security Applications Conference*, pp. 161–170, ACM, 2015.
- [171] "Five pitfalls of cybersecurity insurance: Lessons from the united states." <https://www.lexology.com/library/detail.aspx?g=0afce621-4d25-448e-bb38-1aae0f6c50c7>. Accessed: 2019-07-31.
- [172] "Heads-up: Cyber insurance does not pay out for human error." <https://community.spiceworks.com/topic/1999873-heads-up-cyber-insurance-does-not-pay-out-for-human-error>. Accessed: 2019-07-31.