Modeling Cyber-Insurance: Towards A Unifying Framework

WORKING PAPER*

Rainer Böhme and Galina Schwartz

ICSI and UC Berkeley

Abstract

We propose a comprehensive formal framework to classify all market models of cyber-insurance we are aware of. The framework features a common terminology and deals with the specific properties of cyber-risk in a unified way: interdependent security, correlated risk, and information asymmetries. A survey of existing models, tabulated according to our framework, reveals a discrepancy between informal arguments in favor of cyber-insurance as a tool to align incentives for better network security, and analytical results questioning the viability of a market for cyber-insurance. Using our framework, we show which parameters should be considered and endogenized in future models to close this gap.

1 Introduction

Cyber-insurance, the transfer of financial risk associated with network and computer incidents to a third party, has captured the imagination of professionals and researchers for many years. Yet reality continues to disappoint the proponents of cyber-insurance. Although its roots in the 1980s looked promising, battered by events such as Y2K and 9/11, the market for cyber-insurance failed to thrive and remained in a niche for unusual demands: coverage is tightly limited, and clients include SMBs[1] in need for insurance to qualify for tenders, or community banks too small to hedge the risks of their online banking operations. Even a conservative forecast of 2002, which predicted a global market for cyber-insurance worth $2.5 billion in 2005[2], turned out to be five times higher than the size of the market in 2008 (three years later) [Bae03, BMR09]. Overall, in relative terms, the market for cyber-insurance shrank as the Internet economy grew.

A similar development can be observed in the academic literature. Early works in the 1990s focused on the general merits of cyber-insurance [And94], or protocols

*The authors appreciate comments from readers to keep the survey part of this framework accurate and up to date. The corresponding author can be reached at: rainer.boehme@icsi.berkeley.edu

[1] SMB: small and medium-sized businesses

[2] Conservative, because it is below 2% of total property & casualty premiums; and the forecast was after Y2K, 9/11, and the burst of the Internet bubble.

Workshop on the Economics of Information Security (WEIS), Harvard, June 2010
Last revision: May 21, 2010

borrowed from digital cash to enable risk reallocation in distributed systems [LMN94]. In the late 1990s, when the business perspective of information security became more prominent, visions of cyber-insurance as risk management tool were formulated [Var00, YD02, Grz02, Bae03, GLS03, Sch04b, MYK06, BP07]. These contributions are largely descriptive. If formal, they almost exclusively model the demand side of cyber-insurance (i.e., the trade-off between allocation of security budget on protection mechanisms versus insurance against residual risks). In this literature, the observable underdevelopment of the market for cyber-insurance is often attributed to insurers' lack of experience with a new kind of risk, combined with insufficient actuarial data hindering competitive pricing. Nevertheless, most authors conclude with a positive outlook, in confidence that a resolution of these impediments is merely a matter of time. More recent works acknowledge that the market failed to grow as expected. They attempt to explain market failure with economic equilibrium models, each tailored to one of three obstacles: interdependent security [KH03, OMR05, BL08], correlated risk [Böh05, BK06], and information asymmetries [SSFW09, BMR09]. Their conclusions are more reserved about the prospects of a mature market for cyber-insurance, unless the specific obstacle under investigation could be resolved. However, taking this evidence together, it appears that the market failure can only be overcome if all obstacles are tackled simultaneously. This calls for a comprehensive framework for modeling cyber-risk and cyber-insurance, which also allows us to study the relations between, and the relative importance of, the specific obstacles.

We do not claim to have a silver bullet solution to kick-start the cyber-insurance market, but we have not yet lost our optimism entirely. In this paper, we present a unifying framework which permits to classify the literature and identify areas that have not been covered by the existing models. Our objectives are to take stock, systematize in a common terminology, and give a structured account of a growing field with contributions spread over disperse communities. Our hope is that such a unifying framework helps navigating the literature and stimulates research that results in a more formal basis for policy recommendations involving cyber-risk reallocation [ABCM08, Sect. 9.1]. In addition, we suggest that our framework can be used to partly standardize the exposition of cyber-insurance papers, thus simplifying the tasks of authors' presentation and evaluation of the results by the research community.

One key theme in designing such a framework is to identify factors specific to *cyber*-risk and *cyber*-insurance. This clarifies where novel contributions are needed. Otherwise, one should resort to the standard results for indemnity insurance, which is a well-developed field in economics. However, it largely disregards the specifics of information technology and networked environments.

Our framework breaks the modeling decisions down to five key components: (1) network environment, (2) demand side, (3) supply side, (4) information structure, and (5) organizational environment. Each component covers several model attributes, which imply specific modeling decisions. We discuss all attributes, including their common formalization, with particular emphasis on attributes that are specific to cyber-risk. For less cyber-specific attributes, references to the standard economic literature on indemnity insurance are provided.

This paper makes several contributions. Our proposed framework is presented in Section 2. Within this presentation, the subsection on network environment (Sect. 2.1)

introduces a unified way of dealing with both interdependent security[3] and correlated risk, two obstacles to the development of a cyber-insurance market that so far have been studied only separately. The remaining subsections of Sect. 2 describe the standard economic approach to insurance, augmented to cyber-risk where specific properties arise, in our common notation and terminology. Our terminology is extensible beyond the existing models in the literature to include relevant factors to cyber-risk. These include, for instance, the often-claimed but barely formalized feedback loop to ICT[4] manufacturers, who affect network security via product quality [Böh05, AM06]. Section 3 applies our framework by classifying the relevant literature along the framework's key components. We demonstrate the general usefulness of our framework and its suitability to ease comparisons between different models in a standardized terminology. The framework further permits to pinpoint the driving forces behind the results of models in the literature. Our hope is that this framework will serve as starting point for more systematic extensions in future work by both economists and security engineers. General remarks on the state of the research field and possible directions are discussed in the concluding Section 4.

2 A General Framework for Modeling Cyber-Insurance Markets

Our goal is to develop sufficiently rich framework which unifies the various approaches of modeling cyber-insurance markets in the literature, which is quite diverse. The settings of the existing models differ not only by the particulars of player objectives on demand and supply sides, but also by the assumptions about network structure, player information, actions of the players, and the timing of these actions. To structure this variety, we identified five key components as depicted in Fig. 1.

Our framework includes two natural components, which correspond to demand and supply side of the risk reallocation mechanisms. We make the convention to call parties on the demand side *agents*, and parties on the supply side *insurers*. Most of the specific features of *cyber*-risk are described in a component called *network environment*. In the essence, this component distinguishes a cyber-insurance market from the conventional economic models of insurance. The network environment is composed of atomic elements called *nodes*. Nodes are controlled by agents, who extract utility from the network. This goes along with exposure to risk. We believe the distinction between agents on the demand side and nodes on the network level is useful to separate the business side from the technical risk arrival process. Obviously, in a general framework, we have to allow agents to influence the network environment (see arrow "design" in Fig. 1). The two remaining components are *information structure*, which bundles all modeling decisions that affect the distribution of knowledge among the players about the state of the model, and *organizational environment*, which covers various public and private entities, whose actions affect network security and agents'

[3]Following the economics of security literature, we use the term "interdependent security" to refer to externalities in security decisions. The term does not imply a general reference to statistical dependence, which would subsume correlation.

[4]ICT: information and communication technology

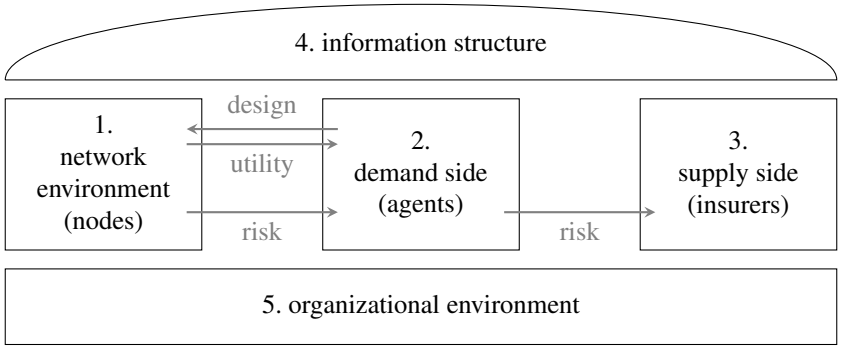

Figure 1: Framework for modeling cyber-insurance markets

security decisions. The latter encapsulates parties who may intervene with the cyber-insurance market although they do not directly appear on the demand or supply side. Awarding information structure a component of its own is justified by the prevalence of information asymmetries in cyber-risk management and their decisiveness in shaping insurance markets in general and cyber-insurance markets in particular. The organizational environment is needed to expand formal models of cyber-insurance markets to a broader system view on cyber-security. We deem such breadth necessary to draw sound and balanced policy conclusions from analytical models.

Before we advance to the details of modeling decisions, let us briefly recall what kind of research questions can be answered with models of cyber-insurance markets. Here, we can distinguish three points of interest:

1. *Breadth of the market*: Looking at the equilibrium condition between demand and supply side, one can pose research questions, such as "Under which conditions will a market for cyber-insurance thrive?" or "What are the reasons behind a market failure and how to overcome them?"

2. *Network security*: Using parameters of the network environment as dependent variable, one can pose research questions, such as "What is the effect of an insurance market on aggregate network security?" or "Will the Internet become more secure if cyber-insurance is broadly adopted?"

3. *Social welfare*: Taking a global perspective, one can account for costs and benefits of all involved parties in welfare analysis and ask questions, such as "What is the contribution of cyber-risk reallocation to social welfare?"

In the following, we will discuss the modeling options for each component.

Notational Conventions We use upper-case letters for functions, $F(\cdot)$; sans-serif letters for random variables, X; and lower-case letters for variables and realizations of random variables, x. Symbols printed bold-face denote vectors, $\boldsymbol{x}$; with components

indexed by a subscript, $\{x_1, x_2, \dots\}$. We slightly abuse the notation by adding subscripts to vectors, $\boldsymbol{x}_{j\neq i}$, which denotes that all components of $\boldsymbol{x}$ except x_i appear as argument of a function. Derivatives are denoted by the operator δ with the argument as suffix, e.g., if $F(x) = x^2$ we have $\delta_x F(x) = 2x$ and $\delta_{xx} F(x) = 2$. This notation is extendable to partial derivatives when needed. Special function $E(\mathsf{X})$ denotes the expected value of X and $P(\mathsf{X} = 1)$ is the probability of the event specified in its argument. Expectations are taken over random realizations of nature in the risk arrival process.

2.1 Network Environment: Connected Nodes

Two properties distinguish cyber-risk from conventional risk. First, nowadays ICT resources are not isolated machines, but interconnected in a network. Their value largely emerges from this interconnection, therefore the analysis of risk and potential losses must take into account the inter-dependencies between connected nodes. Second, most ICT resources are universal automatons and thus have a dual nature: if operational, they generate value for its operators and therefore become loss sources when they malfunction; moreover, when abused or "taken over" by malicious attackers, benign nodes can become threats to other nodes.

In our framework, *nodes* are atomic elements of the network. The risk arrival process is defined at the per-node basis. Note that this bottom-up approach represents the micro perspective. That is, it targets cyber-insurance contract design for an individual agent (who controls a collection of nodes). This corresponds to the *individual risk model* approach in the indemnity insurance literature. Its counterpart is the *compound risk model* [PW92]. In that case, an aggregate perspective is taken that abstracts from micro-particularities of the network. Hence, the latter approach is less suitable for modeling the particularities of cyber-risk occurring on the level of nodes, and we are not aware of any cyber-insurance literature following it.

Recall that our notion of network environment does not necessarily reflect physical connection (e.g., a network link); it includes other forms of interconnectedness, such as logical links or ties in social networks (e.g., for social engineering attacks). Common with the formal literature on cyber-insurance, we abstract from the type of threat: different threats (e.g., targeted attack, viruses and worms, social engineering) may be associated with different network environments [BK06]. In the sense of our framework, real-world cyber-insurance policies covering a range of threats should therefore be understood as a *bundle of contracts*.

We summarize the network environment of models for cyber-insurance by four attributes: defense function, network topology, risk arrival, and attacker model.

2.1.1 Defense Function D

The *defense function* D describes how security investment affects the probability of loss p and the size of the loss l for individual nodes. Its most general form is a probability distribution,

$$p = D(l; s, w, \dots), \tag{1}$$

where s is the level of security[5] and w is the initial wealth, of which the loss l is typically a fraction.[6] Under the assumption of a Bernoulli distribution, a node suffers a potential loss of l with probability p and no loss with probability $1-p$. This simplifies the defense function to,

$$(p, l) = D(s, w). \tag{2}$$

For the Bernoulli distribution, let R be the random variable for losses with realizations $r \in \{0, l\}$. We have $E(\mathsf{R}) = pl$, which is commonly assumed being concave in security investment [GL02], i.e.,

$$\delta_s E(\mathsf{R}) \leq 0,\ \delta_{ss} E(\mathsf{R}) \geq 0. \tag{3}$$

With strict inequalities, Eq. (3) reflects the decreasing returns to investing in security.

The early, simple models allow no security investment and normalize the potential loss l to a constant.[7] Then, p becomes exogenously fixed. To simplify the notation with security investment, let l and w be fixed:

$$p_i = D(\boldsymbol{s}), \tag{4}$$

The agent controlling node i only chooses s_i and takes the vector $\boldsymbol{s}_{j\neq i}$ of all other nodes' level of security as given.[8] This dependence of the defense function for node i on security choices $\boldsymbol{s}_{j\neq i}$ of other nodes is sufficient to model interdependent security [KH03]. If not every node's security choice influences every other node, it is of interest to refine the dependence structure by a model of network topology.

2.1.2 Network Topology G

The *network topology* G is an important attribute to model the particularities of cyber-risk. It describes the relation between elements of an ordered set of nodes. Let function $C_G : \{1, \ldots, n\}^2 \to \{0, 1\}$ be an indicator of connectedness, so that

$$C_G(i, j) = C_G(j, i) = \begin{cases} 0 & \text{if nodes } i \text{ and } j \text{ are connected,} \\ 1 & \text{otherwise.} \end{cases} \tag{5}$$

Technically, G is a lookup function in an undirected unweighted graph, which can be represented in set notation ($C_G(i, j) = 1$ iff an edge exists between vertices i and j), or by a binary adjacency matrix. Extensions to directed graphs or weighted edges are conceivable, but have not been considered in the literature.

Simple models use trivial network topologies, such as

- *independent* nodes,

$$C_G(i, j) = 0\ \forall\ (i, j), \tag{6}$$

[5] To simplify the notation, we view s as a scalar. It would be more exact to view it as a vector, with its components corresponding to the level of security s wrt different threats.

[6] Sometimes the loss exceeds w. For example, this can occur in situations with legal liability.

[7] Constant l is a reasonable simplification to quantify first-party risk. We are not aware of any explicit modeling of third-party cyber-risk, which is more difficult to bound from above.

[8] With the extension of $1 : m$ mappings between agents and nodes in Sect. 2.2.1, agents may actually control sets of associated nodes. Nevertheless, we keep the simpler notation for $1 : 1$ mappings when the difference is not focal.

- or *fully-connected* graphs,

$$C_G(i,j) = 1 \; \forall \; (i,j). \tag{7}$$

Independent nodes represent idiosyncratic risk, and fully-connected graphs allow us to model network externalities in situations where the exact topology appears of minor importance, as can be argued for threats emerging from botnets on globally reachable resources (e.g., certain kinds of phishing or spam). For example, the payout functions for system security as a public good in [Var02] implicitly assume a fully-connected graph because the security level of all nodes is taken into account when calculating the state of any single node irrespective of the aggregation function (total effort, weakest link, best shot).[9] Also the toy examples of two-node cases in [KH03, OMR05, BL08] are (very simple) instances of fully-connected graphs.

To model interdependent security or correlated risk, more expressive topologies are more appropriate. The following types can be found in the literature:

- *star-shaped* graphs underly the single latent-factor model of correlated risk in [Böh05] and the hierarchical treatment of interdependent security in [BL08],

$$C_G(i,j) = \begin{cases} 1 & \text{for } i = 1 \vee j = 1, \\ 0 & \text{otherwise;} \end{cases} \tag{8}$$

- a generalization to *tree-shaped* graphs [LB08b],

$$C_G(i,j) = \begin{cases} 1 & \text{for } q_i = q_j + 1 \wedge C_G(i,k) = 0 \; \forall k : q_k \leq q_i, \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

  with sequence $\boldsymbol{q} = (q_1, q_2, \ldots, q_n)$ containing the cumulative sum of a binary sequence of length n with exactly one leading 0;

- *Erdös-Rényi* (ER) random graphs [LB08b],

$$P(C_G(i,j) = 1) = \text{const} \; \forall (i,j), j \neq i; \tag{10}$$

- *Structured clusters*, the topology behind the two-step risk arrival process in [BK06] to distinguish internal (i.e., within-cluster) from global (i.e., between-cluster) correlation:

$$C_G(i,j) = \begin{cases} 1 & \text{for } q_i = q_j \vee (i-1)(j-1) = 0, \\ 0 & \text{otherwise,} \end{cases} \tag{11}$$

  where $\boldsymbol{q}$ is defined as for Eq. (9) above.

We are not aware of literature using *scale-free* topologies to model interdependent security or correlated risk in cyber-insurance, although this family of graphs is well-established in the areas of reliability modeling and distributed defense for its good fit with real-world networks [AJB00, NA06].

[9] Nodes are called 'agents' in [Var02], as they are not distinguished from players on the demand side of an insurance market.

The network topology can enter market models by shaping the risk arrival process (see the following Sect. 2.1.3), or by defining the information structure when asymmetric information is considered (see Sect. 2.4). However, we are unaware of literature following the latter approach. In principle, layers of multiple network topologies for different properties of cyber-risk are conceivable, e.g., to model the specific influence of social and technical connections [CG08]. Most likely this matches reality better, but unless exact topologies can be collected for real networks and real threats, the additional assumptions will excessively complicate the model.

2.1.3 Risk Arrival

Risk arrival is defined by the relation between the network topology G and the value of the defense function D. We generalize Eq. (4) to

$$p_i = D_i(\boldsymbol{s}, G, [\boldsymbol{x}_{j\neq i}]). \tag{12}$$

As before, p_i for node i depends on s_i, which is chosen by the agent controlling node i, whereas G, like $\boldsymbol{s}_{j\neq i}$, is taken by the agent as given. Vector $\boldsymbol{x} = \{x_i\}$ holds the realizations of a random vector $\mathbf{X} = \{\mathsf{X}_i\}$ introduced to model risk arrival. For simplicity, let $\mathsf{X}_i \sim \{0, 1\}$ be a binary attack state of node i, where the node is compromised if $x_i = 1$, and secure otherwise.

We can distinguish two prototype cases with the defence function D_i being

1. independent of realizations $\boldsymbol{x}_{j\neq i}$ of $\mathbf{X}$, or
2. dependent on the realizations $\boldsymbol{x}_{j\neq i}$ of $\mathbf{X}$.

In both cases, the realized loss r_i at node i depends on x_i, e.g., $r_i = l \cdot x_i$.

Case 1 Let $p_i = P(\mathsf{X}_i = 1)$ be the probability that node i will be compromised, and $\boldsymbol{p} = \{p_i\}$ the corresponding vector of probabilities for all nodes. Therefore, all defense functions D_i depending on elements of $\boldsymbol{p}$ via $\boldsymbol{s}$ (typically the subset of elements where $G_{i,j} = 1$, reflecting topology), belong to Case 1. However, D_i must not depend on any x_j. Thus, in Case 1,

$$p_i = D_i(\boldsymbol{s}, G). \tag{13}$$

Recall that the relation between $\boldsymbol{p}$ and $\boldsymbol{s}$ is given by applying D on all elements of $\boldsymbol{s}$.

Case 2 The dependence of D_i on x_j creates feedback and possibly chain reactions, as p_i stochastically affects x_i and thus feed into D_k of further nodes. We will refer to this property as *risk propagation*,[10] whereas there is no risk propagation in Case 1. Thus, in Case 2,

$$p_i = D_i(\boldsymbol{s}, G, \boldsymbol{x}_{j\neq i}). \tag{14}$$

[10]This property is also referred to as "cascade" in physics or "contagion" in epidemiology. We prefer the term "propagation" borrowed from the reliability and fault tolerance literature (e.g., "error propagation"), since the alternative terms may have unintended connotations in the domain of information security. Moreover, our framework includes features that do not exactly represent the specific notions of cascade (where loss events are deterministic) or contagion (where nodes are assumed to be homogeneous) [LBS09].

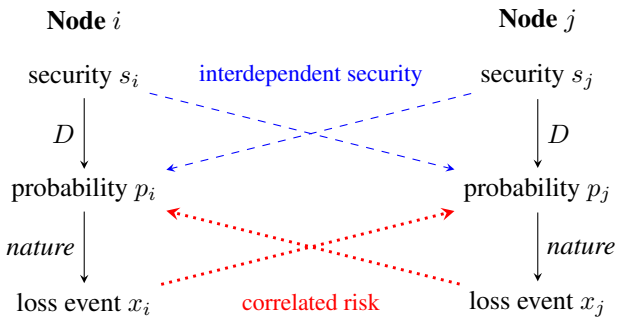

Figure 2: Two-node case illustrating interdependent security and correlated risk in cyber-risk arrival ($G_{i,j} = 1$)

Models with defense functions of Case 1 are easier to tract analytically as simplifying reductions can be found even for non-trivial network topologies, while Case 2 is typically harder. The modeling requires recursive methods or approximations, and may lead to dynamic equilibria [ACY05, GMT05, LBS09].

Observe that our definition of network topology and risk arrival is rich enough to include both interdependent security and correlated risk in a unified way. Indeed, in Case 1, the presence of $\boldsymbol{s}_{j\neq i}$ in the defense function models interdependent security. In the more general Case 2, both interdependent security and correlated risk could be modeled simultaneously via $\boldsymbol{s}_{j\neq i}$ and $\boldsymbol{x}_{j\neq i}$ respectively. Dependence on the realizations $\boldsymbol{x}_{j\neq i}$ is required to model correlated risk (Case 2).

From this analysis, it also becomes apparent that interdependent security and correlated risk have a common root cause: interconnected nodes. The literature on interdependent security focuses on demand-side incentives of this phenomenon, which is then mapped to the supply side via expectations. By contrast, correlated risk is genuinely a supply-side problem. Accordingly, different models of risk arrival are more convenient for the one or the other perspective. Interdependent security can be analyzed on the level of *individual* loss expectations, e.g., $E(\mathsf{R})$ with $\mathsf{R} = l \cdot \mathsf{X}$. This does not require risk propagation. Correlated risk affects risk-averse insurers (see Sect. 2.3) through higher moments of the *compound* loss distribution, e.g., $E(\mathsf{Z}^2)$ with $\mathsf{Z} = \sum_i \mathsf{R}_i$. In networks, higher-order moments of joint distributions are most naturally modeled by risk propagation.

Figure 2 illustrates this by depicting the assumed chain of causality from security level s, which affects the probability of loss p, to the actual realization of a loss event x for two nodes. Interconnection between nodes at various stages of this chain leads to different phenomena, interdependent security or correlated risk. Note that correlated risk is not necessarily a more general case, since it is possible to assume constant security. This removes the problem of free-riding from the model, which is central to typical analyses of interdependent security. By making the one or other simplification, both phenomena were studied separately in the past [Böh05, OMR05]. Although sep-

arate study is possible, we argue that the combination of both is specific to cyber-risk: conventional indemnity insurance considers neither phenomena, examples for interdependent security alone include domains like airline security [KH03], and insurance against natural disasters is a classical example for (locally) correlated risk. Cyber-risk arrival, however, is characterized by both interdependent security and correlated risk.

2.1.4 Attacker Model

The *attacker model* describes actions of assumed perpetrators who intentionally let cyber-risks materialize for their own economic advantage. However, the existing cyber-isnurance literature routinely assumes simple probabilistic rules . This corresponds to a notion of exogenous attacks, being caused by *nature* rather than strategic *players*. Only players' actions are determined endogenously, that would be, attackers react to the agents' and insurers' decisions. Even in the broader field of research on information security investment without explicit risk reallocation, considering truly strategic attackers is uncommon; [LW02, CGK06, CRY08, FG09, Hau09] are commendable exceptions.

Since in reality, a large share of cyber-attacks is strongly strategic, research in this direction appears overdue. With the view of strategic attackers, increased network security may have a positive externality on aggregate network losses, because with higher security, attacker costs might increase, and gains decrease. Attackers reacting to changed incentives might hence seek for alternative (and hopefully more benign) activities.

We suggest modeling attackers as players, with objective functions, information sets, and actions. Our framework naturally extends to include strategic attackers, but it may be hard to choose reasonable assumptions and parameters for their capability. Attackers could be modeled as an additional class of players or as a special type of agents. Note that the tradition in security engineering to anticipate attackers colluding with agents (or insurers) should be maintained in economic modeling.

2.2 Demand Side: Agents

Agents are the entities on the demand side of the cyber-insurance market. They control one or more nodes (e.g., a corporate subnet). By the word "control" we mean that agents make security choices for their nodes and the bear financial consequences when they malfunction. When a cyber-insurance market provides full coverage of these risks, it permits the agents to exchange uncertain future costs with a fixed *premium*, which is paid at present. We use the term "agents" to subsume potential insurance buyers. Firms, consumers or government and non-government organizations can be agents in our terminology. The literature on cyber-insurance is univocal in considering agents as *players* in a game-theoretic sense. The models differ in their assumptions about the agents' endowments and capabilities. In the following we summarize the five relevant modeling options on the demand side: node control, heterogeneity, risk aversion, action space, and time.

2.2.1 Node Control

Node control describes a mapping of each agent to one ($1:1$ mapping) or multiple ($1:m$ mapping) nodes. For simplicity, each node can only be controlled by a single agent. We assume that each agent chooses security investments $\boldsymbol{s}_{i\ldots j}$ for all nodes under his control. His utility depends on the sum of wealth net of losses (if any) of all nodes under his control.

Most models in the literature imply a $1:1$ mapping (and hence may not make the distinction between nodes and agents). However, certain phenomena, such as self-insurance in the presence of two-step correlation [BK06], require a $1:m$ mapping. $1:m$ mappings also appear more realistic for corporate buyers of cyber-insurance who seek coverage for many, possibly geographically distributed, ICT resources. There is no simple generalization from the decision to seek insurance on the level of individual nodes (assuming a $1:1$ mapping) to the same decision of agents who control multiple nodes (i.e., the optimal strategy for agents is not necessarily optimal for each node).

2.2.2 Heterogeneity

Agents (and associated nodes) can be assumed to be homogeneous or heterogeneous in

- l_i, their size of the loss,
- w_i, their wealth,
- D_i, their defense function, and
- U_i, their risk aversion and thus utility function (see Sect. 2.2.3 below).

In case of homogeneity, the respective suffixes can be omitted. We say agents are homogeneous only if all the above-stated properties are identical for them.

Note that a $1:m$ mapping in the node control may imply heterogeneous agents if the network topology is not symmetric for each agent. Further note that homogeneous agents might still face different outcomes through different realizations of their nodes' random variables X_i, i.e., agent homogeneity does not imply perfect correlation.

2.2.3 Agents' Risk Aversion

When the insurance premium is at least actuarially fair, agents seek insurance only if they are *risk averse*, that is, they accept a lower expected value for their income distribution if they can reduce uncertainty. This means reducing the dispersion of the income distribution around its expected value.

Risk aversion is best modeled by a utility function $U : \mathbb{R} \to \mathbb{R}$ which maps monetary wealth w into utility $u = U(w)$. Choosing U concave corresponds to risk-averse agents, i.e.,

$$\delta_w U(w) \leq 0,\ \delta_{ww} U(w) \geq 0, \tag{15}$$

with $E(U(w))$ being the objective function for the agents' optimization problem.

The economic literature distinguishes two commonly imposed utility functions,

- *constant absolute risk aversion* (CARA),

$$U(w) = -\mathrm{e}^{-\sigma w}, \text{ so that } \quad \frac{\delta_{ww} U(w)}{\delta_w U(w)} = \text{const} \quad \text{, and} \tag{16}$$

- *constant relative risk aversion* (CRRA),

$$U(w) = \begin{cases} (1-\sigma)^{-1} w^{1-\sigma} & \text{for} \quad \sigma > 0, \sigma \neq 1, \\ \log(w) & \text{for} \quad \sigma = 1, \end{cases} \tag{17}$$

  so that

$$w \cdot \frac{\delta_{ww} U(w)}{\delta_w U(w)} = \text{const}. \tag{18}$$

In both cases, $\sigma > 0$ is a parameter for the degree of risk aversion. Both forms appear in the cyber-insurance literature. CARA [OMR05] is sometimes imposed for tractability, whereas CRRA [Böh05, MYK06] seems slightly more realistic [Pra64]. Note that risk aversion does not require a behavioral explanation (which could be difficult to maintain in the paradigm of rational firms that leave risk shaping to investors on efficient capital markets). It is sufficient, for instance, to impose a resource constraint for repairing failed nodes to introduce risk aversion indirectly, as shown for a single firm in the context of queuing theory in [CKK05].

2.2.4 Action Space

Established models differ in the action space for agents wrt their insurance purchase. Modeling options are:

- Buying *full* or *partial* cyber-insurance: in case of full insurance, the agent has only binary choice between a full coverage of the potential loss l or no insurance at all, i.e., insurance coverage $\beta \in \{0, 1\}$. Partial insurance lets agents acquire coverage for a fraction of the potential loss, i.e., $\beta \in [0, 1]$.

  Assume for now a $1:1$ mapping between agents and nodes, then the expected wealth for agents who insure against a fraction $\beta_i \in [0, 1]$ of a unit potential loss ($l = 1$) is:

$$E(\mathsf{W}_i) = D_i(\boldsymbol{s}, G)(\beta_i - 1) + w_i - \beta_i \rho. \tag{19}$$

  ρ is the premium for full insurance. In terms of expected utility we obtain:

$$E(\mathsf{U}_i) = D_i(\boldsymbol{s}, G)\, U(w_i + \beta_i - \beta_i \rho - 1) + (1 - D_i(\boldsymbol{s}, G))\, U(w_i - \beta_i \rho). \tag{20}$$

  When the insurer has zero transactional and other expenses, and his expected profit from the contract is zero, then such a contract (and its premium) is called *actuarially fair*. It is a well-known result that at actuarially fair premiums, risk-averse agents strictly prefer full over partial coverage. When partial insurance is available and demanded (i.e., for premiums above the actuarially fair level), it is sometimes difficult to define practical criteria for detecting market failure. This is so because theoretical solutions may exists where agents demand small, yet economically insignificant coverage.

- *Security investment*: agents can *self-protect* by choosing $s_i > 0$ for their nodes and thereby reduce their expected loss (first inequality in Eq. (3)). Taking into account the cost of security investment, we obtain the following expressions for expected wealth (from Eq. (19)),

$$E(\mathsf{W}_i) = D_i(\boldsymbol{s}, G)(\beta_i - 1) + w_i - \beta_i\rho - s_i, \tag{21}$$

and expected utility (from Eq. (20)),

$$\begin{aligned} E(\mathsf{U}_i) = {} & D_i(\boldsymbol{s}, G)\, U(w_i + \beta_i(1-\rho) - s_i - 1) \\ & + (1 - D_i(\boldsymbol{s}, G))\, U(w_i - \beta_i\rho - s_i). \end{aligned} \tag{22}$$

In the standard case of interdependent security, D is defined so that individually rational agents who adjust s_i to maximize Eq. (22) create externalities on other agents' loss distributions [KH03, OMR05].

Some authors allow for a second kind of security investment without externalities and refer to it as *self-insurance*. Self-insurance reduces the size of the loss only for the nodes who invest into it [GCC08, BL08].[11] Self-insurance is mostly studied together with self-protection, so we give the combined expression for the expected wealth:

$$E(\mathsf{W}_i) = D_i(\boldsymbol{s}, G)(\beta_i + \alpha_i - 1) + w_i - \beta_i\rho - s_i - S(\alpha_i), \tag{23}$$

where $\alpha_i \in [0, 1]$ is the level of self-insurance and $S(\alpha)$ is a cost function (typically linear in α). The expected utility can be derived accordingly. To reflect the normative principle of indemnity, which states that insurers only compensate actual damage, a constraint $\alpha + \beta \leq 1$ should be imposed.

- *Endogeneous network formation*: it is conceivable, yet unexplored in the cyber-insurance literature, to consider changes to the network topology as operable actions for agents [GGJ+08]. For example, agents could establish or break up links to other nodes with the intention to reduce their expected loss. We refrain from introducing terminology blindly, but our framework straightly extends in this direction. Such research could bring us closer to the often-stated engineering goal of (re-)designing network architectures in a more insurable way [BL08]. We believe that bottom-up change, driven endogenously by agents' incentives, might be much more implementable than visions of centrally coordinated deployment of a more resilient infrastructure.

  A simple first step would be to consider platform diversity and switching (say, between operating systems) as an endogenous network formation problem. Platforms are represented as disconnected star-shaped segments in the network topology. The center node of each star is the latent factor of a simple risk-arrival model

[11] The term "self-insurance" is borrowed from standard economic literature [EB72], where is has a slightly different meaning. There, it refers to the reduction of the size of the loss (as opposed to self-protection, which reduces the probability of a loss), but neither is connected to externalities. So the term should be used with care, in particular, since the term "self-insurance" is also used without reference to externalities in the cyber-insurance literature [BK06].

with correlation. A node can switch from one platform to another by resolving the connection to the existing center node and establishes a new one to the center node of the target platform.

2.2.5 Time

Simple models of a cyber-insurance market can be formulated for a *single shot*. This means, all choice variables are set only once by all agents (though not necessarily simultaneously, e.g., in [Hof07, LB09, SSFW09], insurers move first, then the agents choose their actions).

When risk propagation is present, even in a single-short formulation, calculating nature's move may require a sequence of updates of D_i as realizations $\boldsymbol{x}$ change. Note that the result may then be sensitive to the initial loss event, i.e., depend on the order in which nodes are updated [LBS09]. To avoid ambiguity, the order should be specified in the model formulation (e.g., single factor models are typically updated from the center of a star-shaped network to its leaves).

Although we cannot foresee any specific benefit that would justify the effort of formulating the market model for *repeated* or *continuous-time* games, initial attempts of dynamic analysis of demand-side security investment come to interesting conclusions, e.g., on the merits of reactive versus proactive security investment [GLL03, BM09]. It remains to be seen if these models, augmented by (simple forms of) cyber-insurance markets, come to genuinely new insights.

As a note of caution, we do not see repeated games as a particularly pressing extension. Repeated games require a static or predictably evolving environment. Rapid technological changes that continue in today's ICT environment, limit the usability of repeated games.

2.3 Supply Side: Insurers

Modeling the supply side with *insurers* as players is essential to analyze market equilibria and the verification of market existence. Early literature has treated cyber-insurance as a choice available to an agent when allocating the security budget. Oftentimes, such papers omit a model of the supply side and hence are "blind" to potential supply-side market failures.

We summarize the supply side by five attributes: market structure, risk aversion, markup, contract design, and higher-order risk transfer.

2.3.1 Market Structure

A central design decision when modeling cyber-insurance is the number of insurers: one (*monopoly*), several (*oligopoly*), or many (*competition*). In the cases of oligopoly and competition, insurers can be modeled *homogeneous* or *heterogeneous*, similar to agents (Sect. 2.2.2).

The dominant model in the cyber-insurance literature is a naive, homogeneous, competitive insurer market structure. Competition is convenient, because it justifies the

assumption that premiums fall to the marginal cost; it is a bit naive though, as the market is assumed to be infinitely large. Partitioning an insurance market reduces the size of each insurers' risk pool, thereby making extreme outcomes more likely. In addition, competition does not always improve the efficiency of insurance markets. Competitive behavior in gaining superior information about agents' risks can even destroy the insurance market [RS76, RS97].

Monopoly, the obvious alternative, is rarely chosen for the difficulty of modeling the demand function, which is an essential input to the insurer's profit-maximization problem.

2.3.2 Insurers' Risk Aversion

A simplification that appears in standard economics textbooks is to assume risk-neutral insurers. However, the insurance industry is regulated in practice to prevent profit-maximizing insurers from taking excessive risk. A typical regulatory measure is a requirement to hold *safety capital*. So it is natural to model risk aversion by this requirement instead of a concave utility function.

Let Z be a random variable of the aggregated loss of all n risks in an insurer's pool for a single period,

$$\mathsf{Z} = \sum_{i=1}^{n} \beta_i \mathsf{R}_i. \tag{24}$$

A risk-neutral insurer breaks even when

$$E(\mathsf{Z}) = \rho \sum_{i=1}^{n} \beta_i, \quad \text{or} \quad E(\mathsf{Z}) = n\rho \text{ for the special case } \beta_i = 1 \; \forall i. \tag{25}$$

However, without additional capital, this insurer would go bankrupt whenever the realization of Z exceeds its expected value; for (approximately) symmetric distributions, this happens in every second period. To prevent this, safety capital c is required so that (for $\beta_i = 1 \; \forall i$)

$$P(\mathsf{Z} > n\rho + c) \leq \varepsilon, \tag{26}$$

where ε is the maximum residual risk of bankruptcy defined by the regulator. Observe that when ε becomes small, the left side of the inequality in Eq. (26) increasingly depends on the right tail of the distribution of Z. Insurers can reduce c if the tail is short, so they prefer lower dispersion of Z. This makes them risk averse (see Sect. 2.2.3).

Due to this dependence on the tail structure, it is relevant to analyze the cumulative distribution of cyber-risk with distribution functions that have parameters for the shape of their tails. Simulation experiments and empirical tests for data on cyber-attacks can be found in [BK06], who model correlated risk with Student-t copulas and estimate from honeynet data, and [MS09], who study the tail structure of data released through breach disclosure laws in the US.

2.3.3 Markup

Let $E(\mathsf{R})$ be an actuarially fair premium. Then,

$$(1 + \lambda) E(\mathsf{R}) \tag{27}$$

is the premiums that corresponds to a *loading* of λ, which can be interpreted in multiple ways:

- *Administrative cost* associated with the contract;
- *Insurer's profit*: a fixed profit extracted per contract, sustainable only for imperfectly competitive and regulated markets;
- *Cost of safety capital*: the cost of holding safety capital $c > 0$ can be distributed on all n contracts belonging to the pool,

$$\lambda = \frac{c \cdot I(\varepsilon)}{n \cdot E(\mathsf{R})}, \tag{28}$$

  where function $I : [0, 1] \to [0, 1]$ gives the real market interest rate demanded for risk with probability of default ε. I is monotonically increasing in its argument. Note that c depends on ε and the distribution of Z (see Eq. (26)).

As correlated risk does not affect the distribution of individual R_i, but the distribution of Z, it affects the existence of an insurance market via λ only for risk-averse insurers. So instead of modeling correlated risk on the level of the network environment, similar outcomes can be obtained (a) by setting the dispersion of Z exogenously and modeling risk-averse insurers, or (b) by imposing a strictly positive λ (e.g., in [OMR05]). Obviously, in both cases, the direct relation to measures of correlation between individual risks disappears.

2.3.4 Contract Design

Contract design by insurers corresponds to and defines part of the agents' action space. The space of contract offers can be modeled as a set of tuples, of which agents can choose elements.

- *Fixed premium*: In the simplest case, the insurer sets a premium ρ for a unit potential loss. Agents can choose β_i.
- *Premium differentiation*: Contracts are tuples (ρ, s) offering premium ρ conditional to a security investment of at least s. To enforce premium differentiation, insurers must be able to observe s. This way, one can model rebates in the insurance premium for better security practices. Assuming symmetric information (for now), premium differentiation can prevent adverse selection and partly internalize the negative externalities of interdependent security [SSFW09]. Extensions of premium differentiation to include self-insurance (ρ, s, α) are conceivable but have not been studied so far.
- *Contract with fines*: Another scantly explored [Hof07, LB09] direction are contracts with fines. Even if s cannot be observed or is too costly to measure at the time of signing the contract, insurers could be in a position to (stochastically) learn about s later on (e.g., when a claim is filed, or just randomly). So agents

and insurers could agree on contracts of the form (ρ, s, f), where f is a fine to be paid by agent i to the insurer in case of a discovery that $s_i < s$. Interesting cases to study emerge if the insurer can observe aggregated properties of $\boldsymbol{s}$, thereby infer the number of 'contract violations', and adjust the intensity of contract monitoring endogenously.

2.3.5 Higher-Order Risk Transfer

Insurers need not be the last step in a chain of risk transfers. Although barely modeled explicitly in the cyber-insurance literature so far, we can distinguish three prototype cases:

- *Cyber-reinsurance*: Most naturally, the idea of reinsurance for dealing with rare catastrophic events seems applicable to cyber-insurance. Modeling reinsurance markets is straight-forward with 'insurers' taking the role of 'agents' and reinsurers become 'insurers'. Instead of $\mathsf{R} > 0$, the loss events become $\mathsf{Z} > \tau$, where τ is a threshold for tail risk. Obviously, reinsurance is more efficient only if reinsurers can pool risks; this assumes the existence of many insurers with *independent* (or at least loosely correlated) risk pools. For conventional insurance branches, this is usually achieved by regional or international diversification. However, due to the global homogeneity of cyber-risk, often attributed to the homogeneity of installed systems [Böh05, GBG+03], cyber-reinsurance is virtually not existent. In January 2002, reinsurers even excluded cyber-risks explicitly from their contracts with insurers in fear of global catastrophic events.

- *Catastrophe bonds*: Modern finance has found countless ways of transferring bundles of risk through financial markets. Catastrophe bonds (short: cat bonds) are financial instruments which pay a decent yield as a risk premium in periods without catastrophic events, but lose their value in case such events occur [D'A92]. Originally developed to facilitate earthquake insurance and related natural perils, cat bonds seem less suitable to transfer cyber-risk. In [ABCM08] it is argued that cat bonds are inadequate for cyber-risks because they may impose adverse incentives on investors, who could improve their financial wealth by causing or commissioning a cyber-attack.

- *Exploit derivatives*: Other financial instruments are tailored more specifically to cyber-risk and avoid such adverse incentives. The concept of exploit derivatives links the payout of the financial instrument to the discovery of vulnerabilities in systems, that is, at a stage before actual losses occur. So even if incentive incompatibilities cannot be ruled out entirely, compared to cat bonds, selfish actions of of individual players are less likely to cause tremendous social damage. Moreover, it is argued that exploit derivatives can form a kind of prediction market [WZ04] to facilitate information sharing about system vulnerabilities, thereby mitigating the information asymmetries prevalent in cyber-security (see also the following Sect. 2.4). However, while exploit derivatives might work for threats related to undiscovered vulnerabilities, this type of threat accounts only for part

of the cyber-risk our society is exposed to. Exploit derivative are informally described and compared to cyber-insurance as risk management tools in [Böh06], but we are not aware of a formal market model including such instruments.

2.4 Information Structure

The importance of information for the provision of network security is generally acknowledged. Considerable research was devoted to developing metrics aiming to provide better information about security levels on various levels of aggregation (agents [Soo00], products [Sch04a], networks [Sha04, GC09]). At the same time, it is theoretically understood and practically observable that strong disincentives keep information sharing below socially optimal levels [GOG05]. Relevant information may not exist, yet it is often the case that it exists but is not available to the decision maker in need for it. This distribution of information is captured in the *information structure*.

Slightly abusing standard game-theoretic terminology, we define *perfect information* as the situation in which no uncertainty is present. Then, no risk is present, and no insurance is needed. Alternatively, information can be *imperfect*, i.e., uncertainty is present. Then, we distinguish two cases: symmetric information and asymmetric information.

We define *symmetric information* as environment with no (i.e., or hidden) private information. Our definition implies that all players, at all times have identical information about the environment, and incur no costs (or delays) associated with information processing.

We define *asymmetric information* as environment where some players have private information, meaning that this information is not available to other players. Situations in which information costs are present and prohibitively high for some players are equivalent to situations where no information exists. Such settings are also covered by our notion of asymmetric information. Observe that it makes no difference at this stage whether the information is factually costly to obtain, or not taken into account for decision making due to agents' cognitive frictions, such as *bounded rationality*.

There exist various interactions between information structure and insurance markets. On the one hand, an exogenously given information structure may dictate the conditions under which an insurance market exists. On the other hand, the presence of cyber-insurers (and an adequate regulatory framework) has the potential to change the information structure for cyber-security. Hence, the information structure has to be treated thoroughly, and ideally endogenously, in market models for cyber-insurance.

2.4.1 Information Asymmetries in the Conventional Insurance Literature

Conventional insurance literature extensively considers the effects of information asymmetries and their ramifications for contract design. Two important problems caused by inferior information of insures about agents are recognized:

- *adverse selection* occurs if insurer cannot distinguish agents of different types ex ante (before the contract is signed),

Table 1: Overview of relevant information asymmetries

	Information	Uninformed party	Informed party
1.	type of nodes	agents	attacker (nature)
2.	agents' choice	insurers	each agent
3.	agents' choice	other agents	each agent
4.	insurers' information	agents	each insurer
5.	insurers' risk pools	other insurers	each insurer
6.	effectivity of protection	agents, insurers	vendors

- *moral hazard* occurs if agents could undertake actions that affect the probability of loss ex post (after the contract is signed and is effective).

Both obstacles are relevant to cyber-insurance, and it has been shown in [SSFW09] that the problem of moral hazard exacerbates when combined with situations of interdependent security. This observation, and the fact that information about security is hard to gather and evaluate (not to mention share it), suggests that pronounced information asymmetries are the third characteristic specific to cyber-risks. To treat this feature adequately, we refine our framework in this direction.

It is known from the economic literature that insurers have two options to structure contracts when they cannot distinguish heterogeneous users: a *pooling* case, where agents of all types are pooled into the same contract; or a *separation* case, where insurers offer two different contracts, and agents sort themselves out according to their type (high or low risks) [RS76].

2.4.2 Information Asymmetries Specific to Cyber-Insurance

Using our framework, we now will identify specific forms of information asymmetries in cyber-insurance. For aspects which have already been touched in the cyber-insurance literature, we will emphasize the link to information asymmetries.

Table 1 shows the most relevant information asymmetries by the content of the information and the uninformed, respectively informed parties. Note that absence of information not only means that this information is not available for decision making. Rather, the existence of *hidden information* available to one party, but knowingly unavailable to another, can be used by the informed party to its advantage. We will briefly discuss rows 1, 2, and 6 of Table 1.

Asymmetric Information: Agents about Nodes (1) Information asymmetries in distinguishing the types of heterogeneous nodes may prevail due to the difficulty of determining the security status of ICT resources. For example, it requires expert knowledge or special tools to identify bots. This is so because compromised nodes do not necessarily exhibit significant performance losses, and certain malware even takes measures to remain undetected. Although not specific to cyber-insurance, [HSA07] discuss

this case of information asymmetry by the example of agents who are unable to distinguish between random failures and intentional attacks.

Related to this is the distinction between *local* or *global* knowledge of agents making decisions for their interconnected nodes. Global knowledge simplifies the analysis of complex network topologies, e.g., by solution concepts such as mean-fields analysis, and has been applied to interdependent security in cyber-insurance [LB08a, BL08]. Aside from methodological tweaks, there exist also functional interpretations: local knowledge appears more realistic a priori, since global knowledge becomes only observable through an intermediary or information aggregator. This is why the availability of global knowledge, and subsequent changes in the equilibrium conditions, can be interpreted as proxy for the effect of information sharing. [GJC09] follow this interpretation and explicitly compare the impact of global and local knowledge for the case of interdependent security.

Asymmetric Information: Insurers about Agents (2) Information asymmetries in distinguishing different types of agents corresponds to the situation described above in Sect. 2.4.1 for conventional insurance. [BMR09] discuss an extension specific to cyber-insurance, where insured agents can opt for off-contract behavior and hide breaches instead of claiming compensation from insurers. They have an incentive to do so if the expected secondary costs exceed the contractually agreed compensation for primary losses.

Asymmetric Information: Agents about Effectivity of Protection (6) Many authors have called a developed cyber-insurance market desirable on the basis of premiums serving as security metrics: price information would counter the lemon problem in the ICT products market [AM06], incentivize vendors to strengthen the security level of their products, and thus improve overall security. Despite these hopes, we are unaware of any direct attempts to model the process of information collection by insurers, nor considering heterogeneous manufacturers as players. Anecdotal evidence suggests that at present, insurers do not engage into collecting and aggregating information about network security. Moreover, it seems to be rather an exception than the norm that cyber-insurers differentiate premiums depending on the security practices of their insured agents [ABCM08].

2.4.3 Timing

The matter of *timing* is related to information asymmetries, important, but not systematically studied in existing cyber-insurance models. Timing involves all modeling decisions on *when* information arrives and is revealed, and to which players. For example, most models assume that agents observe insurers' offered contacts, and then choose security investments, which leads to the conclusion that moral hazard is an obstacle for cyber-insurance. Models with other, equally justifiable, assumptions on timing could come to different conclusions.

2.5 Organizational Environment

Similar to the information structure, attributes of the organizational environment are typically taken as given. However, some arguments suggesting the need for a mature market for cyber-insurance refer to feedback loops with parties outside a narrowly defined risk-transfer mechanism. To study such effects rigorously, parameters of the organizational environment must be included in the models, ideally as endogenous variables.

We have identified four relevant attributes of the organizational environment (sometimes called "stakeholders" in policy contexts): regulator, ICT manufacturers, network intermediaries, and security service providers. In the following, we present options to include each of them into models of cyber-insurance markets.

2.5.1 Regulator

The term *regulator*[12] refers to the government or a governmental authority with power to impose regulation by means of enforceable law or other mandatory rules (e.g., decrees). Hence, the regulator is an important part of the organizational environment for the purpose of policy analysis. A regulator can enter the model in several ways, for example:

- *Disclosure requirements* can improve information for agents and insurers. Here we can distinguish between aggregate and individual agent's information (though in the interest of privacy protection, the use of individual information might be limited in practice).
- *Taxes, fines and subsidies* alter agents' and insurers' costs.
- *Mandatory security impositions* set lower bounds for s. This could be interpreted as introduction of (limited) user liability, because security impositions have to be enforced by the threat of (fixed) fines.
- *Mandatory cyber-insurance* sets lower bounds for β and changes the incentive structure substantially. Also this requirement implicitly introduces a liability regime as the coincidence of connecting a node to the network and obtaining insurance coverage is not natural and needs to be enforceable.
- *Prudential supervision*: The regulator defines the acceptable residual risk ε, the probability of insurer bankruptcy.

Mandatory insurance has been considered as a regulator's tool in [Hof07, BL08]. In the current market environment, such a requirement appears politically inviable and practically difficult to implement [ABCM08]. It seems more suitable to use mandatory cyber-insurance as benchmark case of a social planner, as for instance in [SSFW09]. Liability and fines are discussed in a special section of [OMR05]. We are not aware of specific literature on the other regulatory options in the context of cyber-insurance.

[12]Note that a *regulator* differs from a *social planner* in that the former is a model for a real authority whereas the latter is a hypothetical actor, perfectly informed and omnipotent. They have in common that both seek to maximize welfare. The social planner is usually introduced as benchmark to determine the upper bounds for welfare. In this sense, regulators can be seen as "weaker" social planners.

2.5.2 ICT Manufacturers

ICT manufacturers include vendors of hardware and software equipment. We distinguish two important roles of cyber-insurance:

- *System security*: ICT manufacturers' prioritization of security in the R&D process and their patching strategy affects the defense function D of nodes employing their products. This way, one may conceive a notion of "security productivity", a parameter describing security improvements per unit of security investment for a given technology.
- *System diversity*, notably connected with the market structure, affects correlation in the risk arrival process and thereby the loading λ for risk-averse insurers.

A (formally loose) connection between exogenous market structure and insurance models can be found in [Böh05].

2.5.3 Network Intermediaries

Network intermediaries provision network connectivity services. Internet service providers (ISP) are natural intermediaries; but one can also subsume to this attribute registries, registrars, and application service providers. Possible roles of network intermediaries include:

- contributing to *distributed defense* by sharing information about observed threats or taking down compromised nodes, thereby attenuating contagious risk propagation (the success can be modeled proportional to invested monitoring cost); and
- *shaping the network topology* by establishing or tearing down links strategically. This is related to endogenous network formation (see Sect. 2.2.4).

Obviously, market structure and heterogeneity are relevant factors for network intermediaries, in particular since it is known that the incentives for large and small ISPs diverge due to business-specific factors, such as peering arrangements [ABCM08]. A cyber-insurance market model with ISPs on the demand side can be found in [RKK08], while [GRCC10] investigate intermediaries' role in interdependent security environment. Cyber-insurance in the context of outsourcing [GYL+07] incorporates outsourcing partners as intermediaries. This notion can also be extended to include cloud hosting providers to accommodate this recent trend in the industry.

2.5.4 Security Service Providers

Security service providers include further agents who contribute to network security, e.g., in helping to overcome information asymmetries through collection and aggregation of information, as trusted third parties in information sharing agreements, or to improve information efficiency in monitoring and enforcing contracts (e.g., certifying security levels [PW09]; or conducting forensic investigations when a claim is filed). A

broader notion includes service providers who manage security investment decisions on behalf of the agents. If implementable, this can help to internalize negative externalities of interdependent security [ZXW09].

*

This completes the definition of our framework. In the following we will apply it to survey the literature and identify open research questions.

3 Using the Framework for a Literature Survey

In this section, we demonstrate how our framework can be applied to systematize existing work and identify unexplored aspects. Table 2 shows the relation of components of our framework (in rows) to references (in columns). The table includes only papers which contain (i) technical models of (ii) market-based cyber-insurance. These papers constitute only a small fraction of the entire literature dedicated to the problems of managing cyber-risks. In particular, Table 2 excludes the early papers [Fis02, YD02, GLS03], which do not model the supply side, and the papers focusing on conceptual rather than modeling issues [MYK06, PW09]. Another set of relevant modeling papers [YLG+08, GCC08, Hau09], addressing trade-offs in cyber-risks management or discussing cyber-insurance as a tool without introducing a market in the model, are not included in Table 2 either.[13] Nevertheless, we want to stress the importance of all these works for developing and refining the understanding of the field, including the demand-side attributes of this framework (Sect. 2.2).

3.1 Market Models

Table 2 summarizes existing work on market models according to our framework. We will first present similarities and differences between models and then discuss special features and conclusions of the individual models.

3.1.1 Comparison Across Models

Main Obstacles The framework accounts for three factors that may hinder the development of a market for cyber-insurance: interdependent security, correlated risk, and information asymmetries. In [Böh05, BK06], the focus is on correlated risk, in [OMR05] on interdependent security alone, and in [Hof07, BL08, LB09] on interdependent security with minor information asymmetries. In [BMR09], asymmetric information is studied without features of the network environment. By contrast, [BL08, LB09] solve their model for non-trivial network topologies. Several authors explicitly model the interplay between interdependent security and information asymmetries [RKK08, SSFW09, SSW10]. More precisely, [SSFW09] study moral hazard under different contract monitoring regimes, whereas [RKK08, SSW10] focus on adverse selection.

[13] The term "insurance" even appears in the title of [GCC08], but it refers to self-insurance rather than market insurance. Agents are not assumed to be risk averse in this string of research.

Table 2: Application of the framework to classify cyber-insurance market models

	[RKK08]	[Böh05]	[BK06]	[OMR05]	[Hof07]	[BL08]	[LB09]	[SSFW09]	[BMR09]
Network Environment									
Interdependent security	Y	–	–	Y	Y	Y	Y	Y	–
Correlated risk	–	Y	Y	I	–	–	–	–	–
Non-trivial topology	–	–	Y	–	–	Y	Y	–	–
Agents									
Heterogeneous	Y	Y	–	–	Y	Y	Y	–	Y
Partial insurance	I	–	Y	Y	Y	Y	Y	Y	Y
Security investment	Y	–	–	Y	Y	Y	Y	Y	–
Insurers									
Perfect competition	Y	Y	Y	Y	Y	Y	Y	Y	Y
Risk aversion	–	Y	Y	I	–	–	–	–	–
Markup	Y	–	–	Y	Y	Y	Y	–	Y
Imperfect contracting	–	–	–	Y	Y	Y	Y	Y	Y
Information Structure									
Adverse selection	Y	–	–	–	Y	Y	Y	–	Y
Moral hazard	Y	–	–	Y	Y	Y	Y	Y	–
Organizational Environment									
Regulator	–	–	–	Y[1]	Y[2]	–	Y[1,2]	–	–
ICT manufacturers	–	I	–	–	–	–	–	–	–
Network intermediaries	I	–	–	–	–	–	–	–	–
Security service providers	–	–	–	–	–	–	–	–	–
Research Question									
Breadth of market	–	Y	Y	–	–	–	–	Y	Y
Network security	Y	–	–	Y	Y	Y	Y	Y	–
Social welfare	–	–	–	–	–	–	–	Y	I

Legend: 'Y' = yes, '–' = no, 'I' = informally (but not part of a model), '?' = ambiguous, empty cell = does not apply
Footnotes: [1)] rebates and fines, [2)] mandatory insurance

Demand Side Agents are homogeneous in [OMR05, BK06, SSFW09], and heterogeneous in [RKK08, Böh05, Hof07, BL08, LB09, BMR09, SSW10], often as a prerequisite to study information asymmetries of insurers about agents (see 2.4.1). Contracts with deductibles are standard tools to deal with information asymmetries. Such contracts are introduced in [BL08, SSFW09, BMR09, SSW10]. All models featuring interdependent security, by definition, must allow for some kind of security investment via self-protection. This choice is binary in [Hof07, BL08, LB09] and continuous in the other models. Modeling partial (up to and including full) insurance is common to all models except [Böh05], where only full insurance is considered for simplicity.

Supply Side All models assume homogeneous and perfectly competitive insurers, and most authors include a premium markup. Several authors interpret the markup as a reflection of market power, with higher loading corresponding to more market power [Hof07, BL08, LB09, BMR09]. This interpretation led the researchers to forgo the profit-maximization problem of a monopolist, but their claim of studying the monopolistic market structures is probably too strong. We note that the connection between an exogenous markup and market power is misleading. Alternative interpretations of the markup, such as reflecting a regulated insurance market (see also Sect. 2.3.3), provide equally plausible explanations in stylized models at this level of abstraction. Conversely, if positive markups are observable in practice, this should not be overinterpreted as indicator of market power.

Organizational Environment Current formal models are not particularly rich in capturing parameters of the organizational environment. A link to network intermediaries is established in the informal motivation of [RKK08]. Also largely informally, [Böh05] discusses a comparison of two scenarios in the context of system diversity and the market structure of ICT manufacturers. A formal treatment of mandatory insurance can be found in [Hof07, LB09], who indirectly introduce rebates and fines with this feature. So it remains open if mandatory insurance is a necessary tool, or whether a 'simple' regulation that penalizes agents for underinvesting in self-protection will be sufficient. Rebates and fines are also discussed independently in [OMR05], but the models are not comparable enough to transfer these results to the cases of [Hof07, LB09].

Research Question A single study evaluates its model from the perspective of all three research questions: breadth of the market, network security, and social welfare [SSFW09]. The literature inspired by interdependent security primarily investigates network security, the most natural variable of interest in this setting. By contrast, correlated risk and (ex post) information asymmetries are studied from the point of view of explaining a missing market. The link to social welfare is given in [BMR09] only informally. Strikingly, we are not aware of any study that attempts to capture all three obstacles theoretically and link them with social welfare.

3.1.2 Discussion of Individual Models

Given the above finding on the incompleteness of the theoretical treatment of cyber-insurance, one may ask what the existing results are good for. They clearly give us some intuition on specific aspects and ultimately help to shape a general view, for example the one embodied in this framework. We will now elaborate on individual model features and interpretations using the harmonized terminology of our framework.

The authors of [OMR05] analyze effects of cyber-insurance on incentives for security investment in the presence of interdependent security assuming homogeneous agents and symmetric information. It is shown that with interdependent security, security investments are lower than without. This confirms the seminal result of [KH03] in the presence of market insurance. In a unique analysis, [OMR05] also observe premiums as a function of the competition between insurers, where competition is modeled by declining markups. For stronger interdependence, it turns out that premiums do not necessarily fall because agents shift towards insurance rather than security investment to manage their security risks. Note that due to simplifying assumptions in the proof, the analysis in [OMR05] only holds for relatively small losses compared to the agents' wealth, and for small probabilities of loss.

The analysis in [Hof07] extends this model by heterogeneous agents and asymmetric information. More precisely, security costs are assumed to be distributed uniformly over a continuous interval, and insurers do not know each agents' cost. The result that security investment remains below socially optimal levels is recovered in this setting, and later compared against a situation with mandatory insurance (offered by a single insurer) to internalize the negative externalities of interdependent security.

The arguments in [Böh05, BK06] are exclusively on correlated risk. The author presents conditions under which a market for cyber-insurance is viable despite an assumed monoculture of installed platforms. Specifically, [Böh05] finds that a potential market exists when clients are highly risk averse, and the probability of loss is large. The follow-up work [BK06] refined the argument to two tiers of cyber-risk correlation: *internal correlation*, the correlation of cyber-risks within a firm (i.e., a correlated failure of multiple systems on the internal network), and *global correlation* correlation of cyber-risk at a global level, which also appears in the insurer's portfolio. Also the risk arrival model has been refined from a single-factor Bernoulli model in [Böh05] to beta-binomial models for internal correlation coupled with a Student-t copula on the global level in [BK06]. The authors demonstrate by simulation that a market for cyber-insurance exists for risks with high internal and low global correlation.

In [RKK08], the insurance market is framed in terms of ISPs and their customers. Agents are assumed to have one of two types: high or low probability of loss. Since a pooling equilibrium does not exist with competitive insurers, [RKK08] construct a separation equilibrium (see Sect. 2.4.1). They demonstrate that for some parameters of user types, this separation equilibrium can be destroyed. This way, [RKK08] replicate a famous result in the economic literature [RS76]. If insurers cannot distinguish agent types, there might be no equilibrium. When interdependent security is added to this model, it becomes too complex to derive formal results. So the authors share their intuition about possible equilibrium configurations. In our view, the difficulties encountered in [RKK08] are instructive: with interdependent security, the problem of

contract design becomes intractable very soon (even for two user types). Moreover, as noted in [RKK08], the existence of an equilibrium will be even more questionable than without interdependent security. There seems to be no obvious way out: relaxing interdependent security by introducing sparser network topologies might help the existence of equilibria, but at the same time make the problem even more difficult to tract analytically.

This has been witnessed by a series of publications, [BL08, LB09] (as well as some derivative work not included in Table 2), which features the most comprehensive treatment of network topologies (see also the specific references in Sect. 2.1.2). Agents are mapped to nodes $1:1$ and differ in their defense function so that the cost of security investment varies on a continuous scale. This creates information asymmetries, since insurers know the distribution of security costs in the population, but not each agent's cost of self-protection. In the equilibrium without insurance, only agents with lower security costs invest, thereby replicating qualitatively the result reported in [Hof07] for fully-connected graphs. When self-insurance is allowed as alternative action, complex solution spaces with tipping points emerge, which obstruct a clear interpretation.

A distinct feature of [BMR09] is the assumption of secondary losses which are not covered by insurance contracts nor considered in the premium calculation. This leads to overpriced cyber-insurance products, negatively affecting demand. The paper names reputation losses after security breaches as example for secondary losses and suggests that this kind of overpricing could explain the underdeveloped market for cyber-insurance. The authors adopt a simplified and modified version of [Rav79] to pinpoint the effects of secondary losses and relate this situation to information asymmetry. The paper demonstrates that even without moral hazard and adverse selection, information asymmetry of insurers about agents can lead to overpricing. It remains an open question, however, if contracts can be designed to reduce secondary losses or to separate risks without secondary losses, which then can be priced to match demand.

In [SSFW09] security risks are interdependent, and homogeneous agents can choose how much to invest in security (continuous action). With information asymmetry, that is, when agents' security choices are unobservable by the insurers, the insurance market is missing due to moral hazard. Without asymmetric information, that is, when agents' security investment can be (and is) included into the contract, an equilibrium exists. However, for a substantial parameter range, the possibility to buy cyber-insurance leads to lower security investment. This is so because buying insurance and investing into security are strategic substitutes. When translated to terms of social welfare, [SSFW09] comes to the result common to all models: insurers may have a positive effect on social welfare, but a negative effect on network security. Reallocation of risk helps to reduce security overinvestment and reallocates the freed resources to more productive activities.

3.2 Related Topics

Although our framework was designed to be as comprehensive as possible, there remain a few publications on aspects that are hard to classify. To complete the survey nonetheless, we will discuss them here.

We are aware of literature on two specific sub-problems specific of cyber-insurance

underwriting. First, [IOB09] conducted a qualitative empirical study to identify suitable rating indicators to assess agents' security levels at the time of underwriting. Second, [HH07] suggest a method to translate the number of affected machines to monetary losses by calibrating copula functions.

In [GYL+07, YLG+08], cyber-insurance is proposed to cover risk related to privacy breaches. More specifically, [YLG+08] suggests a random utility model (RUM) to model privacy risks, and an attempt is made to develop insurance instruments using RUM. Although raised as a question in the paper, it includes no analysis of supply-side incentives. A setting with outsourced IT is considered in [GYL+07]. Privacy risks are assumed to be interdependent and informational asymmetric between the agent and the outsourcing provider. These information asymmetries can cause undesirable actions of the provider. [GYL+07] utilizes the principal–agent approach to derive the contract under the assumption of perfectly competitive insurers. The procedure followed by the principla–agent literature is to assume that the principal chooses the contract, or the incentive scheme, to maximize his expected utility subject to constraints assuring that the agent's expected utility is not lower than some pre-specified level (i.e., his incentive constraint holds). [Tir99] describes principal–agent models as "dynamic optimization approach to contracts". While dynamic optimization implies that strategic considerations are only partially addressed, we definitely see some potential in the application of principal–agent literature to cyber-insurance contracts.

However, we remain skeptical about the prospects of privacy insurance given that the known obstacles correlated risk, interdependent security, and information asymmetries are even more acute in the case of privacy, not to mention the 'really hard' practical problems, such as loss substantiation and valuation of privacy [KRG+08, BB09].

4 Concluding Remarks

In this paper, we have proposed a unifying framework for models of the cyber-insurance market. It is unifying in a sense that it unites phenomena that have previously been studied separately, such as interdependent security and correlated risk, in a common risk arrival process. Our framework covers all modeling papers of cyber-insurance markets, and summarizes their results. Despite the early optimism about positive effects of cyber-insurance on network security, by and large, the existing models find that insurance markets might fail. And if a market exists, it tends to have adverse effects on incentives to improve security.

We discussed how the existing literature can be expressed in our unified terminology. This allowed us to compare between the different modeling approaches, to gain insight about which modeling decisions lead to what kind of outcomes. We also identified relevant combinations of properties that are both specific to cyber-risk and not investigated so far.

A common theme of most future research directions suggested throughout this paper is to endogenize parameters that are exogenously given in the existing literature, for instance network topology (Sect. 2.1.2), information structure (Sect. 2.4), and organizational environment (Sect. 2.5). As this clearly affects the analytical tractability, finding the right trade-off between model complexity and expressiveness remains an

important question.

Observe that our framework is limited to analytical models of cyber-insurance. While this reflects the reality of the art—we are not aware of any quantitative empirical work on cyber-insurance markets[14]—, one more oddity of the subject becomes apparent: economic research on conventional insurance followed the existing business practices, whereas researchers of cyber-insurance develop theory in the hope to find a practical business solution.

As a final observation, our stock-taking exercise also revealed a substantial discrepancy between informal conjectures and claims in the cyber-insurance literature on the one hand, and model assumptions and inferences on the other hand. For example, researchers write about how insurers will ...

- ... improve information about security levels. But they do not include parameters that reflect such information improvements in their models.
- ... affect agents' choices of network products (hardware, software, configuration). But existing models of contracts do not reflect these choices.
- ... aggregate information about security (obtained from claims). But they do not model it parametrically.

This list can be continued. It strikes us that many of the positive expectations about cyber-insurance have not been analyzed rigorously, so that our conclusion remain based on weak evidence even after a decade of research. We can only speculate about the reasons; perhaps initial results of market failure have scared researchers' attention off the field, despite considerable interest of policy makers. On the upside, this means that cyber-insurance continues to be an attractive area of interdisciplinary research. Our hope is that this framework contributes to proceed more efficiently.

Acknowledgements

The authors are grateful to Jens Grossklags for sharing his views on an earlier version of this working paper. The first author was supported by a postdoctoral fellowship of the German Academic Exchange Service (DAAD). This work was also supported in part by TRUST (Team for Research in Ubiquitous Secure Technology), which receives support from the US National Science Foundation (NSF award number CCF-0424422).

[14]The empirical part of [BK06] seeks for evidence of correlation in cyber-risk arrival; it does not include market data, such as prices, volumes, or losses.

References

[ABCM08] Ross J. Anderson, Rainer Böhme, Richard Clayton, and Tyler W. Moore. Security economics and the Internal Market. Study commissioned by ENISA, 2008.

[ACY05] James Aspnes, Kevin Chang, and Aleksandr Yampolskiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. In Proceedings of ACM-SIAM Symposium on Discrete Algorithms, pages 43–52, Vancouver, BC, Canada, 2005. ACM Press.

[AJB00] R. Albert, H. Jeong, and A. L. Barabási. Error and attack tolerance of complex networks. Nature, 406:387–482, 2000.

[AM06] Ross J. Anderson and Tyler W. Moore. The economics of information security. Science, 314:610–613, 26 October 2006.

[And94] Ross J. Anderson. Liability and computer security: Nine principles. In Dieter Gollmann, editor, Computer Security (ESORICS '94), LNCS 875, pages 231–245, Berlin Heidelberg, 1994. Springer-Verlag.

[Bae03] Walter S. Baer. Rewarding IT security in the marketplace. Contemporary Security Policy, April 2003.

[BB09] Stefan Berthold and Rainer Böhme. Valuating privacy with option pricing theory. In Workshop on the Economics of Information Security (WEIS), University College London, UK, 2009.

[BK06] Rainer Böhme and Gaurav Kataria. Models and measures for correlation in cyber-insurance. In Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK, 2006.

[BL08] J.C. Bolot and M. Lelarge. A new perspective on internet security using insurance. In Proceedings of IEEE INFOCOM, pages 1948–1956, April 2008.

[BM09] Rainer Böhme and Tyler W. Moore. The iterated weakest link: A model of adaptive security investment. In Workshop on the Economics of Information Security (WEIS), University College London, UK, 2009.

[BMR09] Tridib Bandyopadhyay, Vijay S. Mookerjee, and Ram C. Rao. Why it managers don't go for cyber-insurance products. Communications of the ACM, 52(11):68–73, 2009.

[Böh05] Rainer Böhme. Cyber-insurance revisited. In Workshop on the Economics of Information Security (WEIS), Harvard University, Cambridge, MA, 2005.

[Böh06] Rainer Böhme. A comparison of market approaches to software vulnerability disclosure. In Günter Müller, editor, Emerging Trends in Information and Communication Security, LNCS 3995, pages 298–311, Berlin Heidelberg, 2006. Springer-Verlag.

[BP07] Walter S. Baer and Andrew Parkinson. Cyberinsurance in it security management. IEEE Security and Privacy, 5(3):50–56, 2007.

[CG08] Hsing Kenny Cheng and Hong Guo. Computer virus propagation in a network organization: The interplay between social and technological networks. NET Institute Working Paper #08-24, 2008. `http://ideas.repec.org/p/net/wpaper/0824.html`.

[CGK06] Michael Collins, Carrie Gates, and Gaurav Kataria. A model for opportunistic network exploits: The case of P2P worms. In Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK, 2006.

[CKK05] Pei-Yu Chen, Gaurav Kataria, and Ramayya Krishnan. Software diversity for information security. In Workshop on the Economics of Information Security (WEIS), Harvard University, Cambridge, MA, 2005.

[CRY08] Huseyin Cavusoglu, Srinivasan Raghunathan, and Wei T. Yue. Decision-theoretic and game-theoretic approaches to it security investment. Journal of Management Information Systems, 25(2):281–304, 2008.

[D'A92] Stephen P. D'Arcy. Catastrophe futures: A better hedge for insurers. Journal of Risk and Insurance, 59(4):575–601, 1992.

[EB72] Isaac Ehrlich and Gary S. Becker. Market insurance, self-insurance, and self-protection. Journal of Political Economy, 80(4):623–648, 1972.

[FG09] Neal Fultz and Jens Grossklags. Blue versus red: Towards a model of distributed security attacks. In Roger Dingledine and Philippe Golle, editors, Financial Cryptography, LNCS 5628, pages 167–183, Berlin Heidelberg, 2009. Springer.

[Fis02] Mike Fisk. Causes and remedies for social acceptance of network insecurity. In Workshop on the Economics of Information Security (WEIS), University of California Berkeley, USA, 2002.

[GBG+03] Daniel E. Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quarterman, and Bruce Schneier. CyberInsecurity – The cost of monopoly, 2003. `http://www.ccianet.org/papers/cyberinsecurity.pdf`.

[GC09] Daniel E. Geer and Daniel G. Conway. Hard data is good to find. IEEE Security & Privacy, 10(2):86–87, 2009.

[GCC08] Jens Grossklags, Nicolas Christin, and John Chuang. Secure or insure? A game-theoretic analysis of information security games. In Proceeding of the International Conference on World Wide Web (WWW), pages 209–218, Beijing, China, 2008. ACM Press.

[GGJ+08] Andrea Galeotti, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. Economics Working Papers ECO2008/07, European University Institute, 2008.

[GJC09] Jens Grossklags, Benjamin Johnson, and Nicolas Christin. When information improves information security. Technical report, Carnegie Mellon CyLab, 2009. `http://people.ischool.berkeley.edu/~johnsonb/Welcome_files/paper.pdf`.

[GL02] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. ACM Transactions on Information and System Security, 5(4):438–457, 2002.

[GLL03] Lawrence A. Gordon, Martin P. Loeb, and William Lucyshyn. Information security expenditures and real options: A wait-and-see approach. Computer Security Journal, 14(2):1–7, 2003.

[GLS03] L. A. Gordon, M. Loeb, and T. Sohail. A framework for using insurance for cyber-risk management. Communications of the ACM, 46(3):81–85, 2003.

[GMT05] A. Ganesh, L. Massoulié, and D. Towsley. The effect of network topology on the spread of epidemics. In Proceedings of IEEE INFOCOM, pages 1455–1466, Miami, FL, 2005. IEEE Press.

[GOG05] Esther Gal-Or and Anindya Ghose. The economic incentives for sharing security information. Information Systems Research, 16(2):186–208, 2005.

[GRCC10] Jens Grossklags, Svetlana Radosavac, Alvaro A. Cárdenas, and John Chuang. Nudge: Intermediaries' role in interdependent network security. In Proceedings of ACM Symposium on Applied Computing (INFSEC Track). ACM Press, March 2010.

[Grz02] Torsten Grzebiela. Insurability of electronic commerce risks. In H. Sprague R. editor, Proceedings of the Hawaii International Conference on System Sciences. IEEE Press, 2002.

[GYL+07] S. Gritzalis, A. Yannacopoulos, C. Lambrinoudakis, P. Hatzopoulos, and S. Katsikas. A probabilistic model for optimal insurance contracts against security risks and privacy violation in IT outsourcing environments. International Journal of Information Security, 6(4):197–211, 2007.

[Hau09] Kjell Hausken. Strategic defense and attack of complex networks. International Journal of Performability Engineering, 5(1):13–30, 2009.

[HH07] Hemantha S. Herath and Tejaswini C. Herath. Cyber-insurance: Copula pricing framework and implications for risk management. In Workshop on the Economics of Information Security (WEIS), Carnegie Mellon University, Pittsburgh, PA, 2007.

[Hof07] Annette Hofmann. Internalizing externalities of loss prevention through insurance monopoly: An analysis of interdependent risks. Geneva Risk and Insurance Review, 32(1):91–111, 2007.

[HSA07] Peter Honeyman, Galina Schwartz, and Ari Van Assche. Interdependence of reliability and security. In Workshop on the Economics of Information Security (WEIS), Carnegie Mellon University, Pittsburgh, PA, 2007.

[IOB09] Frank Innerhofer-Oberperfler and Ruth Breu. Potential rating indicators for cyberinsurance: An exploratory qualitative study. In Workshop on the Economics of Information Security (WEIS), University College London, UK, 2009.

[KH03] Howard Kunreuther and Geoffrey Heal. Interdependent security. Journal of Risk and Uncertainty, 26(2-3):231–49, March-May 2003.

[KRG+08] Douglas J. Kelly, Richard A. Raines, Michael R. Grimaila, Rusty O. Baldwin, and Barry E. Mullins. A survey of state-of-the-art in anonymity metrics. In Proceedings of the ACM Workshop on Network Data Anonymization (NDA), pages 31–40, New York, 2008. ACM Press.

[LB08a] Marc Lelarge and Jean Bolot. A local mean field analysis of security investment in networks. In Proceeding of NetEcon, ACM SIGCOMM Workshop, August 2008.

[LB08b] Marc Lelarge and Jean Bolot. Network externalities and the deployment of security features and protocols in the Internet. In Proceedings of ACM SIGMETRICS, pages 37–48. ACM Press, June 2008.

[LB09] Mark Lelarge and Jean Bolot. Economic incentives to increase security in the internet: The case for insurance. In Proceeding of IEEE INFOCOM, pages 1494–1502. IEEE Press, 2009.

[LBS09] Jan Lorenz, Stefano Battiston, and Frank Schweitzer. Systemic risk in a unifying framework for cascading processes on networks. European Physical Journal B, 71(4):441–460, 2009.

[LMN94] Charlie Lai, Gennady Medvinsky, and B. Clifford Neuman. Endorsements, licensing, and insurance for distributed system services. In Proceedings of ACM CCS, pages 170–175, Fairfax, VA, 1994. ACM Press.

[LW02] Kong-wei Lye and Jeannette Wing. Game strategies in network security. Technical Report CMU-CS-02-136, Carnegie Mellon University, 2002. http://www.cs.cmu.edu/~wing/publications/CMU-CS-02-136.pdf.

[MS09] T. Maillart and D. Sornette. Heavy-tailed distribution of cyber-risks, 2009. http://arxiv.org/abs/0803.2256.

[MYK06] R. P. Majuca, W. Yurcik, and J. P. Kesan. The evolution of cyberinsurance. Technical Report CR/0601020, ACM Computing Research Repository, 2006.

[NA06] Shishir Nagaraja and Ross J. Anderson. The topology of covert conflict. In Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK, 2006.

[OMR05] Hulisi Ogut, Ninrup Menon, and S. Raghunathan. Cyber insurance and it security investment: Impact of interdependent risk. In Workshop on the Economics of Information Security (WEIS), Harvard University, Cambridge, MA, 2005.

[Pra64] John W. Pratt. Risk aversion in the small and in the large. Econometrica, 32:122–136, 1964.

[PW92] Harry H. Panjer and Gordon E. Willmot. Insurance Risk Models. Society of Actuaries, Schaumburg, IL, 1992.

[PW09] Manoj Parameswaran and Andrew B. Whinston. Incentive mechanisms for internet security. In H. Raghav Rao and Shambhu Upadhyaya, editors, Handbooks in Information Systems, volume 4, pages 101–139. Emerald, 2009.

[Rav79] Artur Raviv. The design of an optimal insurance policy. American Economic Review, 69(1):84–96, March 1979.

[RKK08] Svetlana Radosavac, James Kempf, and Ulas C. Kozat. Using insurance to increase internet security. In Proceedings of ACM NetEcon'08, pages 43–48, Seattle, WA, 2008. ACM Press.

[RS76] Michael Rothschild and Joseph E. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. The Quarterly Journal of Economics, 90(4):630–49, November 1976.

[RS97] Michael Rothschild and Joseph E. Stiglitz. Competition and insurance twenty years later. The Geneva Papers on Risk and Insurance, 22:73–79, 1997.

[Sch04a] Stuart E. Schechter. Computer security strength and risk: A quantitative approach. PhD thesis, Harvard University, Cambridge, MA, USA, 2004.

[Sch04b] Bruce Schneier. Hacking the business climate for network security. IEEE Computer, 37(4):87–89, 2004.

[Sha04] Shadowserver Foundation. http://www.shadowserver.org/, 2004.

[Soo00] Kevin Soohoo. How much is enough? A risk-management approach to computer security. PhD thesis, Stanford University, 2000.

[SSFW09] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. Competitive Cyber-Insurance and Internet Security. In Workshop on Economics of Information Security 2009, University College London, England, June 2009.

[SSW10] Galina Schwartz, Nikhil Shetty, and Jean Walrand. Cyber-Insurance: Missing Market Driven by User Heterogeneity. In preparation, `www.eecs.berkeley.edu/~nikhils/SecTypes.pdf`, 2010.

[Tir99] Jean Tirole. Incomplete contracts: Where do we stand? Econometrica, 67(4):741–782, 1999.

[Var00] Hal R. Varian. Managing online security risks. New York Times, June 1st, 2000. `http://www.nytimes.com/library/financial/columns/060100econ-scene.html`.

[Var02] Hal R. Varian. System reliability and free riding. In Workshop on the Economics of Information Security (WEIS), University of California, Berkeley, 2002.

[WZ04] Justin Wolfers and Eric Zitzewitz. Prediction markets. Journal of Economic Perspectives, 18(2):107–126, 2004.

[YD02] William Yurcik and David Doss. Cyberinsurance: A market solution to the internet security market failure. In Workshop on Economics and Information Security (WEIS), University of California Berkeley, CA, 2002.

[YLG+08] A. N. Yannacopoulos, C. Lambrinoudakis, S. Gritzalis, S. Z. Xanthopoulos, and S. N. Katsikas. Modeling privacy insurance contracts and their utilization in risk management for ICT firms. In S. Jajodia and J. Lopez, editors, Computer Security (ESORICS '08), LNCS 5283, pages 207–222, Berlin Heidelberg, 2008. Springer-Verlag.

[ZXW09] Xia Zhao, Ling Xue, and Andrew B. Whinston. Managing interdependent information security risks: A study of cyberinsurance, managed security service and risk pooling. In Proceedings of International Conference on Information Systems (ICIS), Phoenix, Arizona, 2009.

List of Symbols

α	fraction of potential loss mitigated by self-insurance
β	fraction of potential loss covered by (market) insurance
ε	risk of insurer bankruptcy
λ	markup over actuarially fair premium
ρ	premium for unit potential loss $l = 1$
σ	parameter of risk aversion
τ	threshold for tail risk in reinsurance scenario
c	safety capital
C_G	indicator function of connectedness between two nodes in G
D	defense function, returns loss distribution (or parameters thereof) as a function of security investment
f	fine for contract violation
G	network topology (see also C_G)
i	primary index for nodes or agents
I	function for market interest rate at given default risk
j	third index for further nodes or agents
j	ternary index for other nodes or agents
l	potential loss, often normalized to 1
m	number of nodes controlled by a specific agent
n	number of nodes, number of contracts in insurer's portfolio
p	probability of loss l
$\boldsymbol{q}$	sequence for formalizing tree-shaped graph and structured clusters
R	random vector of individual losses, realizations $\boldsymbol{r}$
s_i	level of security at node i, expressed in monetary terms of the security investment
$\boldsymbol{s}$	level of security at all nodes
S	cost function for self-insurance
u	utility, random variable U
U	utility function
w	(initial) wealth, random variable W
X	random vector of individual risk arrival, realizations $\boldsymbol{x}$
Z	random vector of aggregated losses