Technische Universität München
TUM School of Computation, Information and Technology

Cyber Risk and Insurance: Risk and Dependence Modelling and Optimal Pricing of Cyber Assistance

Gabriela Angela Zeller

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung einer

Doktorin der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Silke Rolles

Prüfer*innen der Dissertation:

1. Prof. Dr. Matthias Scherer
2. Prof. Dr. Marie Kratz
3. Prof. Dr. Martin Eling

Die Dissertation wurde am 27.06.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 03.11.2023 angenommen.

Technische Universität München

School of Computation, Information and Technology

Cyber Risk and Insurance

Risk and Dependence Modelling and Optimal Pricing of Cyber Assistance

Dissertation

by

Gabriela Angela Zeller

Abstract

This publication-based dissertation is concerned with several topics from the area of cyber risk and cyber insurance. First, a holistic actuarial model for cyber risk based on marked point processes is developed. The model takes into account multiple stylized facts of cyber losses and puts particular emphasis on capturing dependence between cyber incidents in a realistic way. The versatility and practicality of the model for actuarial applications is illustrated in an extensive simulation study. Second, a simplified version of the proposed model is used to illustrate the importance of holistic cyber insurance value chains, in particular the need for sophisticated data collection for actuarial model calibration. Using suitable measures of dependence and risk, we derive analytically that otherwise risk management decisions may be based on a detrimental underestimation of risk. Third, the optimal pricing of novel cyber insurance products including risk mitigation is investigated using the framework of distortion risk measures and a sequential optimization game. Besides representing a new rigorous mathematical approach to pricing cyber assistance services currently offered on the market, the study encompasses a first extension to risk mitigation services which may enhance cyber resilience on a portfolio level.

Zusammenfassung

Diese publikationsbasierte Doktorarbeit beschäftigt sich mit verschiedenen Themen aus dem Bereich Cyber-Risiko und Cyber-Versicherung. Zuerst wird ein ganzheitliches Modell für Cyber-Risiken basierend auf markierten Punktprozessen entwickelt. Das Modell berücksichtigt mehrere statistische Charakteristika und legt besonderen Wert auf die realitätsnahe Modellierung von Abhängigkeit zwischen Cyber-Vorfällen. Die vielseitige Anwendbarkeit des Modells für praktische aktuarielle Fragestellungen wird in einer ausführlichen Simulationsstudie veranschaulicht. Des Weiteren wird eine vereinfachte Version des vorgeschlagenen Modells verwendet, um die Bedeutung der gesamtheitlichen Ausgestaltung von Wertschöpfungsketten in der Cyber-Versicherung, insbesondere die Notwendigkeit der durchdachten Datensammlung zur Kalibrierung aktuarieller Modelle, zu verdeutlichen. Die ansonsten potentiell schädlichen Auswirkungen für das Risikomanagement werden anhand geeigneter Abhängigkeits- und Risiko-Maße analytisch hergeleitet. Außerdem wird die optimale Bepreisung neuartiger Cyber-Versicherungsprodukte, welche Dienstleistungen zur Risikoreduktion enthalten, unter Verwendung von *distortion risk measures* und eines sequenziellen Optimierungsspiels untersucht. Diese Studie stellt einen theoretisch-mathematischen Ansatz zur Bepreisung aktuell am Markt angebotener Cyber Assistance-Dienstleistungen dar und enthält außerdem eine Erweiterung auf Dienstleistungen zur Risikoreduktion, welche dazu dienen können, Cyber-Resilienz auf systemischer Ebene zu verbessern.

Acknowledgements

First and foremost, I am indefinitely grateful to my supervisor, Matthias Scherer, for his continuous guidance, support, and patience. Throughout my doctoral studies, I could always rely on him to be approachable and supportive regarding any questions and challenges, mathematical or otherwise, enabling me to persevere through this marathon project.

My particular gratitude furthermore goes to Rudi Zagst, my academic mentor and supervisor during my Master's studies, whose guidance and advice encouraged me to embark on the journey of pursuing a doctorate.

I would also like to mention the great atmosphere in our research group M13, which allows doctoral candidates to learn from each other without ever feeling competitive. I would like to thank my colleagues for the interesting and often entertaining exchanges during many lunch and coffee breaks and for their great friendship and support, not least for being willing to help out in any circumstances (e.g. by entering my house in hazmat suits!). In particular, I would like to express my deep gratitude to Lexuri Fernández for being an all-round wonderful friend and person, to Henrik Sloot for his constant willingness to share his profound technical and general knowledge, to Yevhen Havrylenko for being as reliable and diligent a colleague as anyone might wish for, and to Janina Engel for her empathy and advice. I would like to thank Bettina Haas for her tireless support with administrative matters and her contribution to a pleasant work environment.

I gratefully acknowledge financial support for my doctoral studies by ERGO Group AG. Moreover, I want to thank all of my ERGO colleagues for welcoming me to their teams, with a special thanks to Konni Kapsocavadis for her empathy and support during a busy time.

I am also very grateful to rely on many great long-time friendships on whose support I can always count and who have continuously encouraged me during my doctoral studies, especially Nadja, Julia, Carina, and all "Kunstakademiker".

Even though they will never read this: I have greatly appreciated the company of our office dog Mila for her steadfast calmness, my cats Whoopi and Pila, whose haughty indifference generally serves to remind one that the importance of many things is only subjective, and Fernando, who constantly made sure I got some fresh air.

Finally, I would like to thank my family who has supported and encouraged my academic and personal development throughout my life, namely my great-aunt Ursula and great-uncle Peter, who have often shared their profound experience and advice, my grandmother Edith and uncle Toni, who could not witness the completion of this thesis but would both have been infinitely proud, and most of all my mother Heidi, whose dedication and perseverance throughout her life have granted me the freedom to pursue my goals and interests and to whom I dedicate this thesis.

List of contributed articles

Core publications as principle author

[1] G. Zeller and M. Scherer. A comprehensive model for cyber risk based on marked point processes and its application to insurance. *European Actuarial Journal*, 12(1):33–85, 2022. https://link.springer.com/article/10.1007/s13385-021-00290-1

[2] G. Zeller and M. Scherer. Risk mitigation services in cyber insurance: optimal contract design and price structure. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 48(2):502-547, 2023. https://link.springer.com/article/10.1057/s41288-023-00289-7

Further articles as principle author

[3] G. Zeller and M. Scherer. Is accumulation risk in cyber systematically underestimated?, 2023. *Working paper submitted for publication, available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4353098.

[4] G. Zeller and M. Scherer. Cyber Insurance: An integral component of Cyber Risk Management. *FIRM e.V. Yearbook 2021*, pages 22–26 and 124–128, 2021. Available at https://www.firm.fm/papers/.

Contents

1 Introduction

Over the last few decades, the digital revolution has moved forward at an ever-increasing pace and information technology (IT) has become the core of many activities, processes, and systems shaping public and private life in today's digitized societies. While the use of smartphones and other connected devices (ranging from personal computers and smartwatches to digitally monitored medical equipment) is ubiquitous in industrialized countries, globally as of April 2023 close to 65% of people were identified as internet users, i.e. belonging to the "digital population" ([130]). People rely on digitized devices for manifold activities in their daily lives, encompassing highly sensitive topics such as private communication and personal data storage. Regarding the corporate sector, some time ago a majority of managers and public decision-makers may have regarded "cyber" as a niche topic to be handled by technical experts within specialized departments. Several large-scale cyber attacks and their public coverage have raised awareness about the dependence of today's business models on functioning IT systems, making their failure or exploitation a substantial business risk. Several studies have investigated the potentially enormous economic cost of cyber crime, which could dwarf the detrimental impact of natural disasters or of the global illegal drug trade (see, e.g., [42]). Alarming as this may already be, the consequences of adverse cyber incidents are unfortunately not limited to purely economic losses, but may include physical damage to property and infrastructure (cf. [14, 13]) and even serious bodily harm to humans (cf. [128]).

Insurers have come to recognize their twofold role in this environment: First, they are themselves large enterprises whose business models integrally depend on the storage and analysis of proprietary and often sensitive data. Second, their core business task of assessing, pricing, underwriting, and managing risks is affected in two ways by the increasing pervasiveness of cyber, namely by implicit exposure to cyber incidents in existing policies (e.g. property or marine insurance) needing to be detected and dealt with as well as – and this will be the focus of this thesis – by the enormous market potential in explicitly insuring businesses against adverse consequences of cyber-related incidents via dedicated cyber insurance policies.

While the topic of cyber risk has received increased research interest by experts from different disciplines (e.g. mathematics, economics, computer science, law), an actuarial modelling approach comprehensively addressing its stylized facts had remained elusive. Its development is a challenging task due to the necessity to balance adequate mathematical complexity with accessibility for actuarial applications in an environment characterized by data scarcity and high non-stationarity. The adequacy of classical actuarial approaches needs to be scrutinized and many actuarial tasks need to be approached in an interdisciplinary way due to the technically intricate nature of information technology.

In the face of the novel and complex nature of cyber risk and its challenging properties, in particular the potential for accumulation risk and heavy-tailed loss severities (see below), insurers tread with increasing caution by limiting coverage and demanding high risk premiums (see, e.g., [123]), often leading to prospective buyers' perception of cyber insurance coverage as currently unattainable (e.g. [102]). This is exacerbated by the fact that an investment in insurance is usually assumed to curtail the limited budget for cyber-security measures, implying the notion of cyber security and cyber insurance as substitutes or even competing activities. Fortunately, many market participants have become aware that – on the contrary – the combination of risk transfer and risk mitigation via so-called *cyber assistance* potentially offers a beneficial enhancement of "traditional" insurance policies for all sides.

This thesis constitutes a contribution to the current research landscape on cyber risk by addressing several of these topics from the viewpoint of insurance mathematics.

The first core article [1] carefully summarizes the existing literature and then develops a comprehensive model for cyber risk using marked point processes based on an extensive review of stylized facts, with a particular emphasis on capturing dependence between cyber incidents and resulting accumulation potential in a realistic way. The model is able to capture the dynamic and heterogeneous nature of cyber risk and is purposely constructed in a modular way to keep it adaptable for actuarial applications. Some properties of particular relevance in an insurance context are derived and the applicability of the model for insurance pricing and risk management is illustrated in a simulation study.

The calibration and testing of the model based on empirical data was not yet feasible due to the scarcity of available data, in particular regarding data which would allow to infer (implicit) dependence information. The main objective of article [3] is thus to transcend the viewpoint of pure actuarial modelling by raising awareness about the importance of designing cyber insurance value chains in a holistic manner. This refers in particular to the often neglected connection between data collection in the claims-settlement department and the statistical foundation needed to calibrate and advance stochastic models in actuarial departments. In [3], we[1] therefore use a simplified (exchangeable) version of the model developed in [1] to illustrate the potentially detrimental effects for risk management if models are calibrated based on claims data where dependence information has been omitted or discarded.

The second core article [2] deals with the optimal pricing of cyber insurance policies including cyber assistance, in particular the question under which circumstances a profit-maximizing insurer would subsidize such services as a way to reduce her overall (portfolio) risk. We extend the study from an interaction between an insurer and a single buyer to a situation in which cyber assistance services can e.g. inhibit the propagation of cyber losses within a portfolio of dependent risks. The results of this study provide a promising starting point for further research about the role of insurers in promoting systemic cyber resilience.

The additional article [4] provides a concise overview of cyber insurance coverage in practice and outlines cyber insurance as one integral part of holistic cyber risk management.

[1] Note that the plural pronoun is used throughout this introduction even though I am the single author of this dissertation. This choice is made for consistency with the contributed articles and – more importantly – to emphasize that the content of this thesis is of course not derived in isolation, but thanks to many opportunities for fruitful discussion with my supervisor and research group colleagues and based on previous work by many researchers of the mathematical scientific community.

2 Actuarial methods & Challenges of cyber

In order to understand why devising and managing cyber insurance products is a major challenge, we briefly explain the foundational ideas and tasks of insurance and their underlying assumptions. These can of course be found in analogous form in any introductory text on the subject (e.g. [122]), and in our view, it is of merit to revise them before studying insurance for a novel risk such as cyber. This is the basis of understanding why certain characteristics of cyber risk challenge the principles of insurance and induce the need to adapt classical actuarial methods.

2.1 Basics of insurance

The fundamental idea

In daily life, every individual or business is exposed to random events with adverse, potentially disastrous, outcomes, for example a lightning strike causing a house or factory to catch fire. Once the possibility of such an adverse event has been identified, one can strive to lower its occurrence probability, e.g. in the present example by choosing appropriate building material and installing a lightning rod. However, as the consequences are so severe, most individuals and companies would still not be willing to invest their life savings or substantial capital into a property that carries the residual (ruinous) risk of burning down at any time (albeit with low probability). From an insurance buyer's viewpoint, an *insurance policy* is a contract with an insurance company that allows to exchange such exposure to a pre-defined, randomly occurring loss event against a fixed upfront payment, called *insurance premium.* Apart from enabling individuals to undertake large endeavours which would otherwise not be feasible due to ruinous risk, such as buying a house in this example, generally an exchange of random losses against certain payments can be assumed beneficial, as a standard assumption about rational decision makers is *risk aversion* (e.g. [138])[2]. The natural question arises how the insurance company in turn benefits from this arrangement.

First, let us emphasize one key advantage of using a specialized insurance company as an intermediary to carry risks of individuals, as opposed to e.g. the historically preceding principle of reciprocal (mutual) insurance, where a large group of individuals would collaborate to share some risk (e.g. inhabitants of a street all saving a little bit of money for the event that one house should burn down).[3] For laymen, random events such as a house fire constitute *uncertainties* in the Knightian sense ([93]), i.e. possible occurrences whose probability cannot be objectively

[2] Risk aversion describes a decision maker's preference for outcomes with a lower rather than higher degree of uncertainty. In other words, the highest payment a decision maker would accept in order to avoid a loss of uncertain size ('certainty equivalent') is larger than the expected size of that loss. The most common way to quantify decision makers' risk aversions is via a utility function, see also Example 1.

[3] Another key advantage is of course facilitating contract closure, akin to the introduction of money to replace pure barter transactions.

quantified. The first task of a specialized insurance company, respectively its mathematicians, lawyers, and technical experts (e.g. civil engineers), is to transform such an *uncertainty* into a *risk* in the Knightian sense, i.e. a measurable quantity that can be assigned "objective"[4] occurrence probabilities. This is accomplished by designing an *insurance policy* that very precisely defines the random events which are covered and the resulting payments dependent on the (random future) manifestation of each event. Additionally, the insurer's experts need to assign probabilities to these random events, but this expert assessment is typically kept proprietary and not shared with the insurance buyer, who only observes it indirectly via the premium which is demanded to insure the resulting risk (see below).

The benefit for the insurance company results from the idea of pooling risks in large homogeneous collectives allowing for *diversification*, both across time and across the collective, due to the law of large numbers. The underlying principle is thus still often described as *"the contribution of the many to the misfortune of the few"*.

Let in the following all random variables (r.v.) be defined on a common, fixed probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For any r.v. X, let $F_X(x) = \mathbb{P}(X \leq x)$ denote its cumulative distribution function (c.d.f.). Let $T > 0$ denote the time horizon (often w.l.o.g. $T = 1$, referring to one policy year). For now, let $n \in \mathbb{N}$ be the number of risks (policies) in an insurance portfolio and let $\{Z_i\}_{i\in\{1,\dots,n\}}$ denote the total annual insurance loss for each risk over one policy year.[5]

Theorem 1 (Diversification in a homogeneous portfolio, see e.g. [115]). *For $n \in \mathbb{N}$, let $\{Z_i\}_{i\in\mathbb{N}}$ be independent and identically distributed (i.i.d.), non-negative, square integrable random variables, and let $S_n := \sum_{i=1}^{n} Z_i$. Then*

$$\lim_{n\to\infty} \frac{\mathbb{V}ar(S_n)}{\mathbb{E}^2[S_n]} = 0, \tag{1}$$

and by the weak law of large numbers, for any $\epsilon > 0$, it holds

$$\lim_{n\to\infty} \mathbb{P}\Big(\Big| \frac{S_n - \mathbb{E}[S_n]}{\mathbb{E}[S_n]} \Big| > \epsilon \Big) = 0. \tag{2}$$

Remark 1 (Notes to Theorem 1).

- *Equation* (1) *is equivalent to stating that the coefficient of variation of S_n tends to* 0 *with growing portfolio size, and means that the standard deviation of the total loss is increasing more slowly than its expected value.*

[4] Of course, as opposed to probabilities for e.g. rolling a fair die, whether assigned probabilities are objective in the strict sense of the word in this case can be debated.

[5] As the r.v.s $\{Z_i\}$ denote insurance losses here, and each policy typically has a *cover limit*, i.e. a maximum amount covered per claim or per policy year, we can assume the distribution of $\{Z_i\}$ to have finite support, implying $\mathbb{E}[|Z_i|^r] < \infty, \forall i, r \in \mathbb{N}$. However, when working in the cyber risk context, it has to be remarked that there is empirical evidence suggesting that the distributions of the underlying risks are quite heavy-tailed and existence of any moments from second order on should not necessarily be assumed, see Subsection 3.1.2 and the references therein. We therefore remark that while we assume (as is customary) square integrability in the following, it is not a necessary condition for the law of large numbers and thus Equations (2) and (3) to hold, see e.g. [80], p. 329.

- *Equation* (2) *states that the probability of the realized loss deviating strongly from the expected loss decreases, relative to the expected loss, as the portfolio size increases.*
- *A sequence of risks is also said to be* "balancing in the collective" *(see [115]) if Equation* (2) *holds. The assumption of independent risks is sufficient, but not necessary, for this convergence, and can be relaxed to uncorrelated risks, and even risks with some degree of positive dependence (see [33] and [115], Theorem 1.6). However, in the latter case, the convergence occurs more slowly, i.e. a relatively larger collective is needed.*

Lemma 1 (Net premium and risk surcharge)**.** *For* $\{Z_i\}_{i\in\mathbb{N}}, S_n$*, as in Theorem 1, the strong law of large numbers holds:*

$$\mathbb{P}\Big(\lim_{n\to\infty}\Big(\frac{S_n}{n}\Big) = \mathbb{E}[Z_1]\Big) = 1. \tag{3}$$

Therefore, for a portfolio of $n \in \mathbb{N}$ *homogeneous risks,* $\mathbb{E}[S_n] = n\mathbb{E}[Z_1]$ *is called* net premium*, as it is the expected capital needed to cover the random total portfolio loss. However, it is also clear by the Central Limit Theorem that*

$$\mathbb{P}(S_n > \mathbb{E}[S_n]) = \mathbb{P}\Big(\frac{\mathbb{E}[S_n] - S_n}{\sqrt{\mathbb{V}ar[S_n]}} < 0\Big) \approx \mathbb{P}(Z < 0) = 0.5,$$

where $Z \sim \mathcal{N}(0,1)$*, i.e. in roughly* 50% *of cases, the net premium would not suffice to cover the random total loss amount, implying insolvency ("ruin") of the insurance company once its initial capital is depleted. A classical result in risk theory (see, e.g., [122]) in the context of the so-called Cramér–Lundberg model (see below), states that ruin occurs almost surely over time for any initial capital* $u > 0$ *if the demanded insurance premium does not exceed the net premium.*

Given the above observation that a *risk surcharge* needs to be added to the expected loss to obtain a feasible risk premium, several common *premium principles* can be considered. For a more exhaustive list and remarks on each principle, we refer e.g. to [33], p. 86f. Note that the reason an insurance buyer would accept a premium that exceeds his expected loss is risk aversion, see Footnote 2.

Example 1 (Some common premium principles)**.** *Let* X *denote an insured risk and* $\Pi(X)$ *the corresponding risk premium for one policy year.*

(P1) Expected value principle: $\Pi(X) = (1+\theta)\mathbb{E}[X]$, $\theta > 0$,

(P2) Standard deviation principle: $\Pi(X) = \mathbb{E}[X] + \theta\sqrt{\mathbb{V}ar(X)}$, $\theta > 0$,

(P3) Zero-utility principle: $\Pi(X)$ *solves* $U(w) = \mathbb{E}[U(w + \Pi(X) - X)]$*, where* w *is the initial capital of the insurance company and* $U(\cdot)$ *is her (strictly increasing and concave) utility function.*

The risk premium that results from such a principle for a specific (collective of) risk(s) X covered by a certain policy is also called *actuarial premium*. The actual premium typically also includes other components based on business considerations, such as an operating expense surcharge, a profit margin, and taxes. Furthermore, as in practice the calculation of the theoretical net premium is of course subject to uncertainty, in particular for new risks, other factors considered for the determination of the premium may be the comparison with similar established lines of business or comparable policies of market competitors. As an extreme case, [121] find *comparison with competitors* and *adaptation from other lines* among the main pricing themes for novel cyber insurance products, hinting at the fact that an established mathematical model to determine an actuarial premium based on a sound understanding of the underlying risk is still elusive.

This is worrisome, as we have seen above that the choice of a premium that exceeds the expected loss does not only imply average profitability[6] of the insurance company, but is a necessary prerequisite for its permanent solvency. Naturally, avoiding a situation of almost sure ruin is a mandatory starting point but by no means sufficient; instead, an insurer must strive to restrict the probability of ruin to an objectively acceptable level. This is - in very simplified terms - the core idea of the *Solvency Capital Requirement* as stipulated in Pillar 1 of the Solvency II Directive ([68]), which constitutes the regulatory backbone of insurer's quantitative risk management in the European Union. Thereby, insurers are required to hold enough capital in order to cover all their obligations (e.g. liabilities from random future losses) over a one-year time horizon in all but 0.5% of (worst) cases (i.e. "ruin" may occur on average at most once every 200 years). These calculations naturally do not consider every individual risk, but are instead based on so-called *risk modules*, which themselves may take into account diversification across lines of business (i.e. are much more complex than a homogeneous collective of risks as above) and are subsequently aggregated to determine an overall capital requirement. Generally speaking, this means an insurer has to determine an adequate amount of capital, denote it $C > 0$, corresponding to the 99.5% quantile of the distribution of a random variable S representing a sum of (unexpected) losses from a group of many (dependent) risks, i.e. based on a chosen stochastic model tackle equations of the type

$$\mathbb{P}(C < S) \leq \alpha \tag{4}$$

where here $\alpha = 0.005$.

In summary, for pricing the insurer is faced with the – by no means easy – task of determining distributions F_X of single risks (or in case of simple premium principles such as $(P1)$ and $(P2)$ at least determining its first moments). Much more complicated still is the task of understanding distributions F_S (i.e. distributions of sums of many (dependent) risks), which is however necessary for quantitative risk management of the portfolio, at the heart of which lie equations of type (4).

[6] This refers to underwriting profit and we do not go into detail about the second source of insurance profit, namely capital investment, as this is of minor importance in non-life insurance. Nevertheless, in principle considering the cost of capital is still relevant here, as it discourages an excessively conservative determination of solvency capital in risk management.

The collective risk model

So far, we have considered sums of a deterministic number $n \in \mathbb{N}$ of total annual losses $\{Z_i\}$. In the non-life insurance context, as each insured risk X can produce multiple claims throughout the policy year, each total annual loss is actually a sum of r.v.s itself, yielding for the annual portfolio loss S_n:

$$S_n = \sum_{i=1}^{n} Z_i = \sum_{i=1}^{n} \sum_{j=1}^{N_i} Y_j^i =: \sum_{k=1}^{N} Y_k, \tag{5}$$

where $\{Z_i\}_{i\in\{1,\dots,n\}}$ is again the annual loss of policy i. Each policy, indexed i, produces a random number of claims N_i with random claim sizes $\{Y_j^i\}_{j\in\{1,\dots,N_i\}}$. If one takes a top-down viewpoint on the portfolio and does not distinguish between policies, the total annual portfolio loss is composed of a random number N of claims of random sizes $\{Y_k\}_{k\in\{1,\dots,N\}}$. This motivates the following general definition.

Definition 1 (Collective risk model)**.** *Consider a portfolio of risks, let $N(t) := (N(t))_{t\geq 0}$ be a counting process describing the total number of claims until time $t > 0$, and let $\{X_i\}_{i\in\mathbb{N}}$, independent of $N(t)$, be an i.i.d. sequence of claim sizes with c.d.f. F_X and $\mathbb{E}[X_1] < \infty$. The total loss (or total claim amount) in a policy year $[0, T]$ is then described by the compound sum*

$$S(T) := \sum_{i=1}^{N(T)} X_i.$$

$(N(T), \{X_i\}_{i\in\mathbb{N}})$ is called ***collective (risk) model****.*

Note that in general the c.d.f. $F_{S(T)}$ cannot be computed in closed form. Thus, in practice, one mostly resorts to simulation or numerical approximations e.g. via *Panjer recursion* ([112]) for distributions of $N(T)$ from the *Panjer class.*

The classical risk model assumes $N(t)$ to be the counting process associated with a homogeneous Poisson process.[7] This implies the number of claims until any fixed time horizon $T > 0$ to follow a Poisson distribution $N(T) \sim \mathrm{Poi}(\lambda T)$, for some intensity $\lambda > 0$. This model, also known as *Cramér–Lundberg model* as it was introduced by Filip Lundberg in [97] and later extensively studied by Harald Cramér ([27]), is the foundation of modern risk theory. [119] emphasizes the importance of the Poisson process in risk theory by stating that the exponential distribution (which is intrinsically linked to the Poisson process via its inter-arrival times) *"plays a similar crucial role in actuarial applications as the Normal distribution does in statistics"*.

While the classical model is theoretically convenient, it relies on simplifying assumptions, e.g. i.i.d. (exponential) claim inter-arrival times. Therefore, since its introduction, many researchers

[7] The intuition for this choice follows quite naturally from the Binomial approximation of the Poisson distribution. If one assumes that the portfolio consists of n i.i.d. risks and the considered time interval $[0, T]$ can be split into independent, smaller intervals such that there can be at most one claim per policy and interval (with probability p), the number of claims for the portfolio follows a Binomial distribution, i.e. $N \sim \mathrm{Bin}(n, p)$. For large n, small p and $\lambda := np$, it follows $N \sim \mathrm{Bin}(n, \lambda/n) \overset{n\to\infty}{\longrightarrow} \mathrm{Poi}(\lambda)$.

have worked on including more sophisticated features in order to include phenomena from reality. Maybe the most natural generalization is to replace the homogeneous Poisson process for claim arrivals by a more general process, e.g. an inhomogeneous Poisson process (deterministic time-dependence of intensity), a general renewal process (*Sparre-Andersen model*, see [16], independent inter-arrival times), a Cox process (stochastic time-dependence of intensity due to exogenous shocks, see e.g. [8]), or a Hawkes process (stochastic time-dependence of intensity due to endogenous shocks, see e.g. [47]). Of course, there are numerous other ways of generalizing the classical model, e.g. by challenging the assumption of i.i.d. claim sizes (e.g. [10]) or the independence between claim arrival times and their sizes (e.g. [9]).

Which risks are insurable?

Reviewing the mathematical foundation of insurance has already alluded to certain prerequisites for a risk to even be *insurable*, i.e. feasible to be covered under an insurance contract. Besides these (implicit) mathematical criteria, there are also economic and societal criteria to be considered. A commonly cited list of insurability criteria is due to [24] and was systematically used in [26] and subsequently addressed in [62] to assess the insurability of cyber risk (see Section 2.2 for the results, and Tables 3 and B2 in [26] for the list of criteria). Below, we briefly summarize some of these criteria (consciously stated rather vaguely), classified as actuarial (A), market (M), and societal (S).[8]

Remark 2 (Some insurability criteria, based on [24])**.**

(A) *Loss occurrences must be independent, random, and allow reliable estimation of loss probabilities.*

(A) *Loss exposures must be such that the maximum possible loss per event is manageable and the average loss per event is moderate. In particular, the expected loss per event must be finite.*

(A) *Information asymmetries (moral hazard and adverse selection) must not be excessive.*

(M) *Insurance premiums must be adequate to ensure sufficient profitability for the insurer while staying affordable for target market participants. Cover limits must be acceptable to both parties.*

(S) *Insurance policies must be consistent with public policy and societal values as well as in accordance with legal restrictions.*

[8] Note that until 2019, most authors emphasized the hopeful perspective that existing challenges to insurability might be alleviated or fully overcome as the cyber insurance market matured. It is true that cyber policies continue to evolve, but unfortunately recent empirical and academic studies show that the cyber insurance market has hardened considerably over the last few years (in particular since the Covid-19 pandemic, which has been labelled by some experts *"the largest-ever cybersecurity threat"* ([104])), i.e. prices have increased, coverage has been limited, and some risks have *"moved towards becoming uninsurable"* (see [102]).

We do not yet comment on or interpret these criteria, but will refer back to them when outlining the characteristics of cyber risk in the next section.

Summary: The tasks of insurance

Having outlined the foundational idea and prerequisites of insurance, we can now summarize the central tasks an insurance company is faced with when devising a new insurance product or line of business.

- **Policy Design:** The risk covered by the insurance policy needs to be clearly defined, in particular which loss events are deemed (un)insurable and included in (resp. excluded from) the insurance coverage. Furthermore, it must be clear which financial consequences of loss events are covered and how their impact can be quantified.
- **Pricing:** Once the insured risk is unambiguously defined, the determination of a feasible insurance premium, fundamentally based on an appropriate stochastic model of the underlying risk to determine the net premium, is a crucial task. As in reality, risks are typically not identically distributed (and Equation (3) can be rather thought of as based on homogeneous subgroups of the whole portfolio), a related important task is **risk assessment**, i.e. the determination of factors along which risks should be distinguished and consequently assigned different premiums. Note that classical pricing focuses on understanding the distribution of individual risks (and therefore relies implicitly on the independence assumption).
- **Portfolio Risk Management:** On the level of the resulting portfolio, the overall risk needs to be measured and quantified, e.g. by monetary risk measures, in order to ensure that the insurance company holds enough regulatory capital to render the probability of insolvency acceptably (vanishingly) small (akin to Equation (4)). Understanding the overall portfolio risk is not only a necessary basis to meet regulatory requirements, but also a valuable assessment of whether the overall risk is still appropriate w.r.t. the risk appetite / strategy of the insurance company, or whether part of the risk should be transferred via the purchase of re-insurance solutions[9]. To this end, it is crucial to understand whether risks in reality can actually be assumed as (approximately) independent, as the extent of diversification potential within a portfolio crucially relies on the absence of positive dependence between risks. Otherwise, it is of paramount importance to understand the factors driving the dependence between risks in order to avoid excessive (relative to the company's risk appetite) accumulation risk in the portfolio.

In the next section, we will outline the challenges the novel risk type "cyber" poses with respect

[9] Throughout this thesis, we consider a primary insurance company's viewpoint, but naturally, the same challenges about cyber risk (accumulation) modelling we tackle are highly relevant from a re-insurer's viewpoint (even with intensified urgency, as re-insurers naturally deal with tail risks).

to the above criteria and tasks, and review the status quo of cyber risk modelling as of the inception of this thesis.

2.2 Challenges of cyber

Over the last decade, cyber has become a topic of great interest for academic researchers, industry professionals, and the general public. Due to ubiquitous digitization and dependence on information technology as well as increasing interconnectivity, reliance on IT and OT (operational technology) systems has become a prerequisite for the functioning of businesses and daily life. As companies from all industry sectors have thus become aware of the exacerbating threat of cyber risk, a rapidly developing cyber insurance market has arisen. For a review of the market and its development, we refer to our summaries in [4] and Section 2.2 of [1] and the references therein, in particular [124, 15, 105] and the excellent survey [100]. Detailed overviews of available cyber coverage are e.g. given in [121], [100] (Table 1), and [26] (Table 2).
For a novel risk type such as cyber, the first challenge is to establish a suitable definition, i.e. to find a consensus about *what even is cyber risk.* The plethora of potential definitions is e.g. alluded to in Appendix D of [61], where 20 different definitions are compared. Clearly, the suitability of such a definition depends on the context, and for an insurance application, it needs to allow for the derivation of a taxonomy of cyber events and their consequences which constitute the basis of an insurance policy. To this end, we find the following definition by [61], in line with the concise definition by [36] in the context of operational risk management, most useful as a basis for our modelling approach, see Section 3.1 of [1].

Definition 2 (Cyber risk, [61])**.**
Cyber risk is "[a]ny risk emerging from the use of information and communication technology (ICT) that compromises the confidentiality, availability, or integrity of data or services. [...] Cyber risk is either caused [naturally] or is man-made, where the latter can emerge from human failure, cyber criminality (e.g. extortion, fraud), cyberwar, and cyber terrorism. It is characterised by interdependencies, potential extreme events, high uncertainty with respect to data and modelling approach, and risk of change."

A well-known comprehensive categorisation of cyber risks is due to [41], but should be understood as a unifying framework or a *"starting point for discussion"* ([41]) rather than a basis for insurance modelling, see also the corresponding assessment in the recent survey [21]. We therefore introduce our own taxonomy of cyber risks, based on the three classical information security protection goals as in [31], see Table 2 in [1]. Definition 2 already alludes to the dynamic and complex nature of cyber risk, and indeed, many empirical studies (see Section 2.2.1 in [1]) have conjectured that the development of the global cyber insurance market has been hampered by its particularly challenging nature and the resulting question to which extent cyber is an insurable risk. A detailed analysis of the latter question, based on empirical data and related literature, was conducted in [26] and subsequently addressed in [62]. The most problematic features identified therein include the lack of independence of loss occurrences, the presence of information asymmetries, and the coverage complexity and its restriction by cover limits. It is interesting to note that these main challenges are not restricted to the actuarial domain, but include economic / market criteria (compare the categories in Remark 2). Likewise, it is worthwhile to remind oneself that these challenges can and should not be tackled from a merely mathematical / statistical viewpoint (e.g. by devising sophisticated stochastic models), but interdisciplinary collaboration with technical and legal experts, e.g. regarding adequate policy design and risk assessment, is essential.

We briefly summarize the key challenging properties of cyber risk. These have been discussed extensively in the rapidly growing body of recent research on the topic, such that we do not give an exhaustive list of references here and refer to Section 3.1 of [1] and the surveys [100, 62, 137] for details.

Remark 3 (Key challenging properties of cyber risk, e.g. [100, 62])**.** *The key challenging properties of cyber risk include:*

- ***Dynamic risk type:*** *Cyber risk is subject to non-stationarity due to the rapid evolution of the threat landscape, the underlying technology and the legal framework, which makes the usability of past data for modelling future losses difficult. Furthermore, many companies' exposure is changing due to increasing dependence on IT systems, compare Remark 5.*
- ***Lack of data:*** *The novelty of cyber risk and the absence of an established terminology for cyber incidents complicate the creation of reliable databases with information on losses. This is exacerbated by reporting bias and the high non-stationarity of the risk.*
- ***Strategic threat actors:*** *Cyber losses do not occur in a completely random fashion, as they are often caused by malicious actors with strategic (economic) motives and attack patterns.*
- ***Interdependence / Accumulation risk:*** *The interconnectedness of IT-systems and the often systemic nature of vulnerabilities induce a dependence structure within and across company networks and the potential for loss accumulation.*
- ***Interdependence of security:*** *The interdependence of loss occurrences and the strategic motivation of threat actors may result in an interdependence of security measures, which in a game-theoretical context might lead to systematic underinvestment in security measures across companies.*
- ***Difficult impact determination and potentially extreme impact:*** *Due to the intangible nature of information assets, it is often difficult to quantify the economic consequences of a cyber incident, in particular completely immaterial impacts such as e.g. reputational damage. Furthermore, some impacts of cyber incidents (such as a long-lasting business interruption) may have ruinous consequences.*
- ***Information asymmetry:*** *Cyber insurance exhibits adverse selection, relating to the*

> *difficulty for an insurer to reliably determine a company's risk exposure, and moral hazard, i.e. the difficulty of ensuring the risk exposure to be maintained throughout the policy year.*

Throughout our work, we exclude the explicit modelling of moral hazard (incomplete information) as we regard it as a minor issue in practice for several reasons: In reality, the main problem is incomplete information about this novel risk for all parties rather than asymmetric information (see also our detailed discussion of cyber risk assessment services in Online Appendix A.1 of [2]). Furthermore, as some consequences of cyber incidents (in particular reputational risk) can clearly not be transferred via insurance, insurance buyers are intrinsically motivated to avoid an incident, even if insurance coverage for other financial consequences is in place. As to the latter, we have alluded to the fact that due to the challenging nature of cyber risk, most insurance policies are endowed with strict cover limits, thus the insurance buyer always carries a residual risk (and a substantial one, assuming that the impacts of cyber losses can be heavy-tailed, see below).
It is nevertheless quite obvious that the other properties in Remark 3 challenge the classical actuarial assumptions in Section 2.1 and the criteria in Remark 2. Particularly the lack of independence of loss occurrences is alarming, as it inhibits the potential for diversification and entails accumulation risk, which is often deemed the most worrisome aspect of cyber risk for insurers (see e.g. [86, 117]). The presence of strategic threat actors challenges the assumption of the randomness of loss occurrences, as they are at least partially caused by malicious agents with their own economic motives. The strong non-stationarity puts into question the assumption of identically distributed risks (even for the same incident type within a short time period), and combined with the lack of reliable data and difficult impact determination complicates the task of determining suitable distributional assumptions for an actuarial model and its subsequent calibration and back-testing. This is exacerbated by the potential presence of extreme cyber loss severities, which complicate the calibration of (standard) models and require additionally the toolbox of *extreme value theory.*

In the last years, many researchers and practitioners have studied cyber risk from their viewpoint. We refrain from providing an exhaustive overview here and refer to our extensive literature review in Section 2 of [1] for the corresponding references. Other excellent reviews are given in [53] (with a focus on business and actuarial science), and earlier [62] (Table A.1) and [61] (Table 5, with a focus on (dependence) modelling approaches).
We identified several main research streams: Earlier works (until 2010) were mostly motivated by a game-theoretic viewpoint on interdependent security on networks, studying questions about (socially optimal) equilibria and their relation to the presence of an insurance market. An in-depth overview of those studies is given e.g. in [100] (Tables 5-7) and [109] (Table 2). More recently, many researchers have tackled the modelling of the interdependence of cyber losses from their provenience, e.g. using copula approaches, (marked) point process models, methods from time series analysis, and models of epidemic spreading on networks originating from mathematical biology (which have of course received increased attention over the last few

years due to their use in the modelling of virus spreads). Furthermore, other works have been dedicated to the statistical investigation of (the few) publicly available databases on cyber incidents / losses and sometimes additionally the collection of new useful cyber data (e.g. by extracting cyber losses from an operational risk database, see [63], or unearthing completely new feasible data sources, see the recent work of [43]). These studies provide invaluable information for the development and calibration of stochastic models, e.g. by hinting at suitable distributional assumptions for the frequency and severity of cyber losses, the covariates and time-dynamics that govern them, and the necessity to distinguish between *"cyber risks of daily life"* and *"extreme cyber risks"* ([63]).

In summary, despite the existence of a variety of recent research approaches to cyber risk, we found that the overall picture remained fragmented and an established modelling approach, bridging the gap between capturing the challenging properties of cyber and staying accessible for widespread insurance application, remained elusive. The need for multidisciplinary research on cyber risk and cyber insurance was emphasized in several excellent surveys on the topic, see [52, 72, 75, 58] and corroborated by the founding of many international initiatives (see e.g. Table 6 in [61] for an overview) and expert working groups, let us mention the German Actuarial Association's (DAV) group on "Data and methods for the valuation of cyber risks" as one example. Among the most relevant actuarial questions for future research identified in the above surveys were the development of a set of actuarial models capturing the properties of cyber risk (in particular the description of extreme cyber scenarios as well as the dependence structure between cyber losses and resulting accumulation scenarios), the identification of factors inhibiting the growth of the global cyber insurance market, and the design of new business solutions to transfer and reduce cyber risks and thus improve the resilience of the overall economy and society.

3 Developing a holistic model for cyber risk

As all of the above studies emphasize, the complex and multi-faceted nature of cyber risk requires a multidisciplinary approach. This need for a multitude of perspectives is explicitly emphasized e.g. in the recent editorial of the Special Issue of the journal *Risks* dedicated to "Cyber Risk and Security" ([44]):

> *[...] We specifically looked for contributions also coming from disciplines other than statistics and actuarial mathematics, to enlarge the perspective and provide distinct, complementary, and sharp insights to researchers of actuarial mathematics or risk analysis and management. It should help improve the understanding of cyber risk, which by nature requires a pluridisciplinary approach if we want to tackle this complex risk in an innovative and relevant way. Research is generally focused on narrow*

> *specialization and people choose to concentrate on highly rated or specialized journals. This Special Issue is an effort to break these barriers and we thank the researchers who have taken the courage to contribute to this unusual special volume. Readers will see how rich those various approaches are and how creative research on cyber risk can be. [...]*

Therefore, our first step before choosing a mathematical modelling approach was to inquire the opinions of several IT security experts and cyber insurance industry professionals[10] about the nature of cyber risk, in particular the mechanisms of the underlying dependencies between cyber incidents and subsequently the characteristics of monetary insurance losses. We then recognized our first principal task as actuarial science researchers in translating this technical knowledge, enriched with findings from the existing academic and empirical literature, into a mathematical model which captures the stylized facts of cyber risk while staying tractable and suitable for practical insurance applications.

3.1 A marked point process model

As described above, many researchers have recently approached cyber risk modelling from their provenience. One approach we deem particularly suitable for actuarial applications is the framework of *marked point processes*, as it provides great flexibility to incorporate relevant stylized statistical facts while yielding actuarial models which can stay conceptually close to the classical frequency-severity-approaches insurers are used to and feel comfortable working with (see Section 2). The latter point is a matter of realism rather than convenience, as in our view the adoption potential of a model in practice (due to operational constraints, acceptance by end-users, and regulatory requirements) is substantially increased if the approach constitutes a suitable "extension" of established and implemented methods rather than a completely new paradigm.

In the following, we briefly detail and put into context two principal stylized facts of cyber risk we aim to capture in our modelling approach, namely the dependence between cyber incidents entailing accumulation risk and the potentially heavy-tailed severity of cyber losses.

3.1.1 Modelling dependent cyber incidents

To illustrate the flexibility of marked point processes, we provide a general definition (see also Online Appendix A.1 of [1] and the references therein, in particular [46] for a general introduction to the topic) and some examples of their application.

Definition 3 (Marked Point Process ([46], 6.4.I and Prop. 6.4.IV))**.** *A marked point process (MPP) with locations in the complete separable metric space (c.s.m.s.) $\mathcal{X}$ and marks in the*

[10] See the Acknowledgements of [1] for details.

c.s.m.s. $\mathcal{K}$ is a point process $\{(x_i, k_i)\}$ on $\mathcal{X} \times \mathcal{K}$ with the additional property that the ground process $N_g(\cdot)$, meaning the process of locations $\{(x_i)\}$, is itself a point process, i.e. for bounded $A \in \mathcal{B}_{\mathcal{X}}$, $N_g(A) = N(A \times \mathcal{K}) < \infty$.

Let N be a MPP with independent marks. Then the probability structure of N is completely defined by the distribution of N_g and the mark kernel $\{F(k|x) : k \in \mathcal{B}_{\mathcal{K}}, x \in \mathcal{X}\}$ representing the conditional distribution of the mark, given location x.

For many applications, one can specify $\mathcal{X} = \mathbb{R}_+$ and consider $\{x_i\}$ as locations along the "time axis" representing the arrival times of events under consideration.[11] In the present context, we will focus on examples where the ground process belongs to the classes of simple point processes mentioned in Section 2.1, referring to the points in time of some event occurrence.

Example 2 (The versatility of marked point process models)**.** *(Marked) Point process models have been suggested for a variety of applications; we briefly provide one classical and two cyber-specific examples.*

- ***ETAS model for earthquake occurrences:***[12] *A marked Hawkes process with $\mathcal{X} = \mathcal{K} = \mathbb{R}$ is used to describe earthquake occurrences in the ETAS (Epidemic Type After-Shock) model. The locations $\{x_i\}$ are occurrence times, where the stochastic intensity of the Hawkes process is governed by* Omori's Law *for earthquake aftershocks, and the marks $\{k_i\}$ are magnitudes, whose distribution corresponds to the Gutenberg–Richter frequency-magnitude law. An important extension is given by letting the ground process include space coordinates, such that the Hawkes branching structure refers to evolution in both space and time.*
- ***Marked Hawkes process model for cyber attack rates ([113]):*** *A marked Hawkes process with $\mathcal{X} = \mathcal{K} = \mathbb{R}$ is suggested for modelling extreme cyber attack rates. The locations, denoted $\{t_i\}$, represent arrival times of extreme attack rates, whose stochastic intensity follows an Autoregressive Conditional Duration (ACD) model to accommodate stylized facts observed in the considered data (e.g. the slow decay of autocorrelation between inter-exceedance times). The marks, denoted $\{\tilde{x}_i\}$, represent threshold exceedance magnitudes and their density thus follows a generalized Pareto distribution, parameterized conditional on the history of the process (i.e. the ground process as well as previous marks).*
- ***Multivariate Hawkes process model for data breach attacks ([25]):*** *A multivariate Hawkes process with $\mathcal{X} = \mathbb{R}$ is suggested to model arrivals of data breach attacks for different groups (e.g. relating to the type of breach, industry sector of the affected entity, or geographical location). The arrivals for each group are determined by a base intensity and a matrix of kernels describing self-excitation within groups (along the diagonal) and*

[11] Naturally, the definition is not restricted to such cases, and the locations $\{x_i\}$ could e.g. more literally refer to two- or three-dimensional location coordinates on the earth's surface representing earthquake occurrences (where the marks could encode the strength of the quake) or something completely abstract altogether.

[12] For more details, see [46], Example 6.4(d) and the original work of [110].

cross-excitation between groups. Among several different kernel specifications of exponential decay, one with non-instantaneous excitation is found to be most suitable for the considered data. If one is not explicitly interested in the form of self- and cross-excitation, one can again consider the overall arrival process as a marked point process, where the marks flag arrivals as belonging to different groups.

In our view, choosing a ground process with such a fairly complicated intensity function is most suitable for applications where the (short-term) temporal branching structure is of principal interest and where data is available in suitable amount and granularity to fit the parameters of the process. This is usually the case in some financial applications (such as limit-order book trading) and when looking at pure cyber attack rates (e.g. traffic measured at a *honeypot*), where multiple occurrences are recorded per second and timestamps are typically available with millisecond granularity. For those applications, methods from the realm of (high-frequency) time-series analysis and the analysis of the suitability of rather complicated arrival processes are certainly of interest (see, e.g., [38, 147, 148]).

In an insurance context, where one is interested mainly in the (average) frequency of loss occurrences over a medium term (such as a policy year) and claims data is collected at most with a daily timestamp, we find it more suitable to remain close to the classical model and assume an inhomogeneous Poisson process as the ground process. Apart from the more realistic possibility of fitting such a model to (prospective) empirical insurance data, the Poisson process class is theoretically opportune due to its many convenient analytical properties, such as being closed under superposition and thinning (see Online Appendix A.1 of [1]). This "easy" choice of ground process allows us to choose a rather complicated two-dimensional mark space encoding two interesting stylized facts of cyber risk:

- The possibility of cyber events to cause dependent incidents in an insurance portfolio, where the probability of insureds being affected jointly by an event depends on underlying common risk factors (such as industry sector affiliation or usage of the same operating system).
- The potential resilience of companies to deter a cyber incident from manifesting itself and causing an actual monetary loss by the establishment of adequate controls, whose effectiveness depends on the sophistication of the attack.

Remark 4 (Network models)**.** *A complementary actuarial research stream on cyber risk is dedicated to models of epidemic spreading on networks (see Section 3 of the recent survey [21] for details and a list of references). These models – originally used in mathematical biology and epidemiology – are quite a natural candidate if one considers a cyber virus as a contagious infection spreading between entities or machines comparable to a biological virus spreading in a population. Phenomena of interest in such models are, e.g., the dynamics of an epidemic within a population, i.e. the state probabilities (to be susceptible to infection, infected, or recovered) of individual nodes at any point in time, or the overall infection time in the system. A particularly* *interesting topic in the cyber context concerns the effect of different network topologies on the before-mentioned quantities. We make the following comments about this class of models:*

- *We refrain from their application in the actuarial context for two reasons: First, according to the expert opinions we collected, the "epidemic spreading" assumption is representative only for a minority of cyber attacks (albeit with potentially very severe consequences), namely so-called worm-type viruses which can indeed propagate between connected machines. In contrast, the majority of connected attacks is due to the potential to exploit a common vulnerability. Second, their quantitative applicability is (as of today) limited by the prohibitive complexity of the data (i.e. extensive underlying network structures) needed to fit them to real-world scenarios. Some authors (e.g. [84, 85]) strive to derive very coarse assumptions from the few, large historical cyber epidemics such as the global WannaCry attacks in 2017 (see, e.g., [70]), but a realistic, quantitative actuarial application seems yet out of reach.*[13]
- *Nevertheless, network models can be considered useful to derive qualitative results, in particular regarding the effect of different network structures on the resilience of the economy / society as a whole with respect to cyber epidemics. Let us in particular mention the ongoing work of [20], who study security- and topology-based interventions to control systemic cyber risks. This refers to the idea of developing resilient networks not only by having individual nodes equip themselves with security investments to be protected from cyber contagion, but by intervening on the level of the network structure, e.g. by edge removal and node splitting. These studies provide starting points for qualitative evaluations of existing regulatory measures e.g. relating to critical supply chains.*
- *Lastly, let us remark that models of epidemic spreading on networks and (marked) point processes are by no means disparate: If one is not necessarily interested in the underlying mechanisms, but rather the times and total number of infections from a top-down perspective, as outlined above a* Hawkes *process is a natural approach to reproduce endogenous clustering (i.e. contagion within a population), see also Section 3.1 of [21]. [84, 85] reasonably argue that while a typical contagion model may be suitable to describe a cyber epidemic in the global population, an insurance portfolio typically only constitutes a small part of this population. Therefore, on the level of the insurance portfolio, rather than dealing with endogenous contagion, during a cyber epidemic one observes a period with increased exogenous contagion. This relates to the idea of using a* Cox *process to describe the evolution of incident numbers in the portfolio.*

Another very popular tool to assess dependence are copula approaches, as by *Sklar's theorem* (see [129]) they allow the decomposition of a multivariate distribution into the marginal distributions of the components and an object, called copula, representing the dependence structure,

[13] Some experts emphasize the hopeful perspective that this may change in the medium to long term due to the potential of corresponding artificial intelligence approaches.

which itself is a multivariate distribution function (with standardized marginals). In the cyber insurance context, some works have aimed to identify copulas to describe the dependence observed in empirical data ([114, 54]), while others have suggested corresponding theoretical approaches ([83, 28]). In particular the first kind of works provide invaluable insights, both about the empirical dependence between cyber losses and by simply raising awareness that concentrating on the statistical analysis of suitable distributions for the marginals it not sufficient. The drawback of copula models is that they constitute pure top-down approaches without aiming at understanding the underlying mechanism of the dependence. In a domain such as cyber plagued by data scarcity and high non-stationary, this is clearly not ideal, as it is likely that there will be an ongoing necessity to combine (historical) statistical evidence with state-of-the-art expert opinions in order to calibrate forward-looking models, a task that is accomplished more easily in a bottom-up model.

3.1.2 Modelling extreme cyber loss severities

Another characteristic of cyber risk that insurers are particularly worried about is the potential of heavy-tailed loss severities of cyber incidents.[14] From a business perspective, cyber risk is usually regarded as part of operational risk, whose heterogeneity, potential heavy-tailedness, and general difficulty of management has been long investigated (see, e.g., [108, 37, 65]). Some authors have therefore approached cyber from this perspective, e.g. by aiming at embedding it into classical operational risk taxonomies ([36, 39]) or by extracting cyber-related entries from an operational risk database and comparing (and finding significant differences between) statistical properties of cyber vs. non-cyber operational risk ([63, 55]). Other authors have analysed empirical data about cyber losses directly looking for potential heavy-tailedness (e.g. [99, 50, 140, 139], see also Table A.1 in [63] for details). Due to scarcity of data about general cyber incidents, these studies have mostly been confined to data breaches based on the well-known publicly available *Chronology of Data Breaches* database by the California-based nonprofit corporation *Privacy Rights Clearinghouse (PRC)*[15], where the severity of an incident is measured as the number of affected records.[16] Exceptions are [63, 39] and the recent work by [43] who investigate a new interesting database of cyber complaints filed with the French police (for details on the dataset, see Section 3.1 therein) they obtained for their investigation. An overview of the results of these studies w.r.t. to the fatness of the tail distribution, is given in Table 1 of [45].

[14] Note that by referring to cyber incidents instead of cyber events here, we refer to cyber-related loss severities experienced by a single policyholder. Taking a macro-economic perspective, aggregate losses from cyber *events* can of course also exhibit heavy-tailedness as multiple entities are affected simultaneously or in short succession by a cyber incident. However, in an actuarial context, we find it more straightforward to subsume the phenomenon of simultaneous, dependent incidents under the umbrella of frequency modelling (see previous subsection) and consider severity modelling of cyber *incidents* only.

[15] Available for public download from `https://privacyrights.org/data-breaches`.

[16] The most cited approaches for a subsequent conversion to an approximate monetary loss are *Jacob's formula* ([90]), which was however based on a very specific sample and its general applicability has been cautioned against by the author himself, and a subsequent amendment by [73], see e.g. Online Appendix A.4 of [1] for details.

Following [63] and akin to [43], we suggest a combination of lognormal and Generalized Pareto distribution (GPD) for the severity of general cyber incidents. The choice of a GPD for the tail distribution is by no means arbitrary, as it is rooted in the framework of extreme value theory (EVT), which provides the appropriate theoretical toolbox to investigate extreme events and heavy-tailed data. For a general introduction, we recommend e.g. [94] or the excellent introductory book by [64]. In brief summary, there are two main pillars of univariate EVT: On the one hand, the so-called *three-types theorem* due to Fisher–Tippett ([74]) and Gnedenko ([79]) provides a result on the asymptotic distribution of the renormalized maxima of a sample (akin to the Central Limit Theorem for renormalized means). As the name suggests, there are three possible candidate distributions for a non-degenerate limit (which can be unified into a single parametric family, called *Generalized Extreme Value* distribution), corresponding to the baseline distribution being short-tailed, light-tailed, and heavy-tailed. On the other hand, the Pickands–Balkema–de Haan theorem shows that when considering *extreme* observations above a sufficiently high threshold, the excess distribution function (i.e. the distribution of the exceedances above that threshold) can be approximated by the Generalized Pareto distribution, where again there are three types of tails (short, light, and heavy) depending on the sign of the tail index (or shape parameter[17]) ξ.[18] The density of the GPD with shape parameter ξ and scale parameter β is given by

$$GPD_{\xi,\beta}(x) = \begin{cases} 1 - \left(1 + \frac{\xi}{\beta}x\right)^{-1/\xi}, & \text{if } \xi \neq 0, \\ 1 - \exp\left(-\frac{x}{\beta}\right), & \text{if } \xi = 0, \end{cases}$$

for $x \geq 0$ if $\xi \geq 0$ and $x \in [0, -\beta/\xi]$ if $\xi < 0$. This is the basis of the so-called *peak-over-threshold* (POT) approach, which aims at fitting a GPD to understand the tail distribution of a data sample (in the cyber context, see e.g. [63, 73, 43] and the overview Table 1 of [45]) and illustrates the ubiquity of this choice of distribution to model tail behaviour.[19]

The immediate question which arises for the application of the POT approach is the appropriate choice of threshold to balance the accuracy of the GPD approximation and the size of the underlying sample (i.e. the higher the threshold, the more theoretically accurate is the approximation, but the fewer data points lie above it and can thus be used for the estimation). Various methods for this task have been suggested, mostly relying on graphical examinations of the so-called mean-excess and Hill plots (see, e.g., [118], p.85ff.) or on rule-of-thumb methods like simply taking the highest $x\%$ of observations. [63] choose a bootstrap goodness-of-fit test suggested by [132]. For some unsupervised methods, see [48] and the references therein. A very promising iterative algorithm to choose a threshold "automatically" (i.e. in an unsupervised fashion) has been developed and recently applied to cyber data by [43]. The method provides

[17] Note that some authors use both terms synonymously, while others denote as *shape parameter* $\frac{1}{\xi}$.

[18] The cases are $\xi > 0$ corresponding to a heavy tail (Fréchet maximum domain of attraction), $\xi = 0$ to a light tail (Gumbel maximum domain of attraction), and $\xi < 0$ to a short tail (Weibull maximum domain of attraction). In the cyber risk context, the first case is of interest.

[19] Note that the existence of such a theoretical framework is particularly invaluable because it conceptually allows to draw implications about extreme events that exceed any previously observed occurrences in the sample.

additional flexibility by suggesting up to three components, namely a lognormal distribution for the body and a GPD for the tail, joined by an exponential bridge which automatically vanishes during calibration if appropriate from the data. In fact, this automatic removal of the exponential bridge by the iterative fitting procedure is observed for the considered cyber database, yielding a two-component (lognormal-GPD) model which had been previously suggested for another cyber data set by [63] and thus as a general approach for the severity distribution of cyber incidents in [1]. [63] furthermore employed the approach of [37] for time- and covariate-dependent fitting of the GPD to investigate the dependence of the tail distribution parameters on characteristics of the affected company. While this is an important step to assess the heterogeneity of cyber losses, let us emphasize that all approaches so far choose a global threshold for the POT approach. In our view, an interesting task for future research would be a procedure to determine a covariate-dependent threshold, e.g. by developing a suitable EM-type procedure to fit a mixture model (of body and tail distribution). This corresponds to the idea that depending on the covariates (e.g. for companies of varying size), different observations should be considered *extreme*.

Remark 5 (Heterogeneity of cyber risk)**.** *At this point, let us briefly interject a remark on the heterogeneity of cyber risk, which is another major challenging factor and can be understood in three main dimensions:*

- ***Heterogeneity of threats****: As alluded to in Definition 2, cyber risk subsumes a variety of categories w.r.t. origins, causes, and consequences, whose potentially different statistical nature (see the references in Section 3.2.1 of [1]) needs to be understood and incorporated into a model. The relevance of certain categories for a particular actuarial model depends on the chosen taxonomy and the underlying insurance product. For classifications of cyber threats, we refer to the sources in Section 2.2 and the references therein.*
- ***Heterogeneity of impact****: Another dimension concerns the heterogeneity of targets of a cyber incident w.r.t. their exposure and resulting potential negative consequences. It is crucial for an insurer's risk assessment to understand how characteristics of a prospective policyholder affect the frequency and severity of a certain cyber incident. For empirical analyses on the influence of certain covariates (e.g. company size), we refer to the references in Section 3.2 of [1], in particular [63], and the recent works [56, 55].*
- ***Heterogeneity over time****: The above-mentioned non-stationarity can be regarded as a third kind of heterogeneity, namely referring to the changing nature of both previous points over time. Various empirical studies have addressed the question of time trends in cyber loss data*[20]*, coming to conflicting conclusions (most likely due to a large variety in underlying data and applied methodologies), see e.g. [59] and the references therein.*

As to the first point, we suggest a classification of cyber risks along two axes, which allows to

[20] This refers to studies on non-stationarity of both frequency and severity, as well as studies investigating all kinds of cyber incidents or "extreme" cyber incidents only (i.e. investigating stationarity of the distribution or tail distribution).

incorporate accumulation risk and distinguish different types of cyber incidents in a way which aims to balance comprehensiveness and applicability in an actuarial context, see Table 2 in [1]. Considering the second and third point, we suggest a time- and covariate-dependent modelling approach, based on previous empirical findings from the literature, which can be fit to available data with established statistical methods.
With respect to all the above possibilities to incorporate heterogeneity, we suggest a modular modelling approach with the purpose of rendering it easily adaptable for any end-user's application, i.e. to amend or replace parts as dictated e.g. by updated future insights on the nature of cyber risk or properties of an internal data set, without having to replace the overall structure.

The aforementioned works studying the tail properties and potential extreme occurrence of cyber loss severities are particularly relevant in an insurance context, as the potential for diversification effects in a portfolio may be thwarted by the presence of very heavy-tailed distributions (see, e.g., [60, 89, 88]). This concern is of course exacerbated when coupled with the existence of non-linear dependencies as in the cyber context (see previous subsection).[21]

[45] compare the tail index fits of several studies to corroborate that most evidence points towards the finiteness of the first moment (as a necessary criterion for insurability) to be fulfilled for cyber risk. At the same time, [43] emphasize that with a tail index indicating potential non-existence of any moments higher than first order, cyber should be classified as a very high risk comparable to natural catastrophes.[22]

Remark 6 (Insurance cover limits)**.** *Note that for actuarial modelling, finiteness of moments for individual severity distributions is not a technical concern, as the size of the insurance claim corresponding to a cyber incident is contractually equipped with a cover limit, i.e. the distribution of the insurance loss is truncated at the cover limit and therefore technically short-tailed. Nevertheless, this does of course not diminish the need to strive to understand the characteristics of the underlying risk which is then mapped to a claim size distribution. Apart from the insurer's necessity to assess the tail of the portfolio loss distribution to accurately determine regulatory risk capital, let us emphasize that cover limits are a contractual feature designed to limit the insurer's liability in rare cases and not expected to actually come into force regularly. Furthermore, as insurance buyers consequentially carry the residual risk exceeding the insurance cover limit, they may rightly cast doubt on the usefulness of insurance policies with relatively small cover limits for heavy-tailed risks. There are convincing arguments for the claim that this phenomenon actually affects the real cyber insurance market: While insurers are treading with*

[21] Note that while we consider dependence purely as part of frequency modelling, e.g. the first approach suggested in [19], based on [116], advocates for modelling dependent severities.

[22] While indeed due to its statistical properties, in particular the heavy-tailedness of loss severities, cyber is often likened to *NatCat*, we have emphasized one key difference in Remark 3, namely the man-made, strategic character of cyber as opposed to the purely random (governed by stochastic physical laws of nature) occurrence of natural catastrophes. Another difference concerns the potential for diversification, which for the latter quite obviously corresponds to geographical considerations, whereas it is less clear along which lines to diversify for cyber insurance portfolios.

increasing caution and therefore equip cyber policies with relatively strict cover limits, prospective insurance buyers are looking for coverage precisely in case of extreme scenarios like a large data breach or long-lasting business interruption. The perceived non-existence of adequate cover limits was cited as one reason for companies to refrain from closing a cyber insurance contract in [6, 7].

Two conclusions can be drawn: First, as emphasized by [45], a task of paramount importance concerns the need for the collection and analysis of more representative data on cyber incidents in order to better understand the statistical properties of this novel risk type. While the complexity of this task is increased by the heterogeneity and high non-stationarity of the cyber risk landscape, empirical data analysis is an elementary prerequisite for the conceptualization, evaluation, and further development of any proposed stochastic model. Second, for a risk type like cyber, where the application of traditional actuarial approaches and the design of suitable insurance policies with an acceptable premium to both sides are challenging – compare the remark about the offer-demand-mismatch regarding cover limits above – novel insurance solutions transcending mere risk transfer need to be devised as a way to close the market gap and increase cyber resilience.
We will discuss these topics in the following Sections 3.2 and 4, respectively.

3.2 Scarcity of data or scarcity of information?

As described above, we identified as first task for an actuary to understand the underlying characteristics of cyber risk and to translate them into a mathematical model which is flexible, accessible to industry professionals, and adaptable to the peculiarities of real-world insurance portfolios. At many points, concrete choices within the modelling structure were dictated by the paradigm of imposing as little structure as possible (e.g. Poisson process for frequency modelling, Uniform attack strength distribution, and industry sector as single (stylized) underlying factor for common vulnerabilities) given lack of evidence from suitable data that a more complicated choice was warranted. As an initial step to illustrate (qualitative) results of the model, we relied on expert opinions and findings from empirical studies to choose parameters for an exemplary simulation study.
Naturally, several additional sensitivity analyses could have been conducted, and more generally speaking, many interesting stochastic models for cyber risk could be devised. However, while we believe these endeavors to be theoretically worthwhile for mathematical researchers, they remain somewhat academic exercises. While numerous academic works on cyber risk have been published in recent years, the gap between "research" and true practical "innovation" in the cyber insurance domain has persisted, if not widened. Therefore, from our vantage point at the intersection of academia and practice, a currently urgent task is the evaluation (and subsequent further development) of the existing proposed modelling approaches for cyber based on suitable data rather than the conception of additional theoretical models.

Data scarcity is ubiquitously cited as one of the major challenges in the cyber risk context. It is of course true that representative and reliable data on cyber losses is scarce. This is the reason that those works dedicated to analyzing the scant available public databases and unearthing or creating new databases provide particularly invaluable insights about the statistical nature of cyber risk (e.g. [63, 140, 59, 43]). However, the comparison of their (sometimes diverging) results is complicated by the heterogeneity of the underlying data and their representativeness for a particular insurance application is by no means guaranteed.
Therefore, in the realm of cyber insurance, both in theory and practice, the development, fitting, and back-testing of models is still severely hampered by the lack of suitable available data. Currently, researchers are struggling to collect any available data and make it fit for insurance applications, e.g. as described above by mapping data breach "severities" to economic costs and then approximating insurance losses.

Naturally, the most reliable and representative data source for an insurer's actuarial modelling purposes is their own historical claims database, with two obvious limitations: First, cyber insurance portfolios, due to their short time in existence and comparative small size as well as the low frequency of cyber losses, cannot be expected to generate an abundant amount of claims data compared to long-established, larger insurance lines like automotive. Second, due to the high non-stationarity of the cyber risk landscape, the validity of historical claims data should unfortunately be expected to decline comparatively quickly.

Nevertheless, there is another problem which we encountered when considering suitable databases to fit our model: In the (cyber) insurance value chain, actuaries' tasks of stochastic modelling for pricing and risk management are often handled separately from related tasks like risk assessment and product design, but in particular separately from claims settlement, which is usually conducted by a completely disjoint group of (legal) experts, who do not have actuarial modelling aspects on their agenda of primary concerns, see Figure 1 in [3]. The crucial connection between data collected in the claims-settlement department and data needed to fit and backtest a model in actuarial departments is usually overlooked. This leads to data being collected "naïvely", i.e. following established lines of business, individual information about each claim is stored in a way that (ideally) allows to infer individual policyholders' loss distributions. However, crucial information about the cause of the loss, which would allow to identify dependent claims stemming from the same event, is sometimes available in unstructured form (e.g. text within IT forensic reports which are stored for each claim separately), but never connected with or added to the structured claims database. In order to illustrate the potentially detrimental effects of such missing information on portfolio risk measurement, in [3] we use a purposely simplified (exchangeable) version of the model suggested in [1] and compare results for such a model with full dependence information to a model based on (partially) missing information. This illustrates the urgent practical necessity of creating holistic cyber underwriting processes and establishing meaningful data collection.

We definitely agree that data collection is one of the current tasks of paramount importance in order to develop and refine cyber insurance models and thus improve the understanding of this challenging risk; “data collection” is rightly identified as one major research topic in [137]. We emphasize, however, that the task of gathering data has to be approached in a way which allows to perspectively gain relevant *information* about the risk - crucially, in the cyber insurance context, this does not only refer to data about marginal distributions, but also about dependence between losses. We therefore see it fit to caution insurers against “naïve” data collection following established paradigms from existing lines (where independence may rationally be assumed and the estimation of marginal distributions is sufficient), which may have them end up with more raw data on cyber losses in a few years' time, but still no (or only partial) information about the full extent of the underlying risk.

4 Cyber assistance: Novel insurance products as a way forward

In the last sections, we have outlined the actuarial challenges regarding the development of cyber insurance products and have hinted at the fact that the real-world market, despite its rapid growth being regularly predicted, has lagged behind the expected development and has recently hardened with respect to prices and available coverage, see, e.g., [102, 81, 101]. All parties are interested in ways to close the persistent market gap, and one potential way to approach this challenge is the conception of novel insurance products containing risk mitigation measures alongside classical risk transfer, so-called *cyber assistance*. In the face of cyber as a potentially systemic risk whose insurability is regularly questioned, we find this idea particularly promising and have therefore studied the optimal pricing of such combined insurance products in [2]. While the suggestion of risk mitigation services on an individual contract level already exists in the market (see Section 4.1), we would already like to emphasize a further advancement of this idea, namely the hopeful perspective that insurers may make positive use of the dependence between cyber losses in the portfolio:[23] By using their knowledge about observed losses, they may be able to warn other policyholders who are particularly vulnerable to an imminent cyber incident and therefore prevent future losses, thereby contributing to the resilience of the whole system (see Section 4.3).

4.1 Risk management strategies

We have so far focused on the question of quantifying cyber risk for a classical actuarial application. This relates to insurance as a means of *risk transfer* as introduced in Section 2.1: A prospective insurance buyer, who identifies his business as being exposed to an unacceptable extent of potential adversity from incidents related to cyber, approaches an insurer to have this

[23] Note that so far, positive dependence has been introduced as a purely dangerous, worrisome feature of cyber risk, as it hampers diversification and potentially entails accumulation risk.

quantified as a risk within an insurance policy and to rid himself of the random exposure for a corresponding deterministic premium. Naturally, this is not the only type of action a company could take with respect to cyber risk; in [4], we have outlined the consideration of cyber insurance as one component of a holistic cyber risk management strategy. In general, after a risk has been identified and assessed, there are different classical ways of coping with it.

Remark 7 (Four ways of coping with risk)**.** *Traditionally, there are four possibilities of dealing with an identified risk (see, e.g., [100]), namely*

- *risk acceptance,*
- *risk avoidance,*
- *risk mitigation,*
- *risk transfer.*

It is clear that due to the ubiquity of digital technology and the increasing dependence of all parts of business on information technology, the first two options are not feasible for cyber.[24] *Risk transfer via cyber insurance is receiving increased attention by companies across sectors and sizes (see, e.g., [107]), but its uptake continues to be limited, whereby high prices, unavailability of desired coverage, and policy complexity remain the main obstacles (e.g. [107]).*

It seems reasonable to state that some level of cyber risk mitigation exists at every company, ranging from elementary security measures (like anti-virus software and password-protection of devices) at small businesses to extensive protection and incident response plans, accounted for and coordinated by whole departments of IT security experts, at large enterprises. Over the last decade, the perception of cyber security has shifted from being viewed as a merely technical topic within IT teams and systems to being considered a substantial business risk which needs to be incorporated into enterprise risk management (ERM). This heightened awareness of the dependence of modern business models on information technology is likely the main driver of companies' focus on cyber risk mitigation, accompanied by customer expectations, concerns about reputation risk in case of a cyber incident, and the increase of legal and regulatory requirements in many jurisdictions.[25]

In theory and practice, risk mitigation and risk transfer are often regarded as separate and even competing activities which need to be paid from the same limited budget, and therefore e.g. an

[24] Some years ago, a non-negligible portion of small businesses may have been content to adopt the strategy of risk acceptance, assuming fallaciously that their exposure was negligible. This view has likely changed over the last few years due to widespread media coverage of cyber incidents, the shift to remote work during the Covid-19 pandemic, and generally heightened awareness of ever-increasing reliance on digital technologies (see, e.g., [77] for the development of the interest in cyber insurance among small and medium-sized enterprises (SMEs) in Germany).

[25] In the European Union, this ranges e.g. from the General Data Protection Regulation ([66], in force since 2018) affecting all businesses to the extensive Digital Operational Resilience Act ([69], entering into force in 2025) for the financial sector, to name only the tip of the iceberg.

investment in cyber insurance implies reduced spending for cyber security measures. Some game-theoretical studies on interdependent security on networks indeed conclude that the presence of insurance inhibits investments in protection measures in the system (e.g. [126, 125]). Such conclusions are often exacerbated by worries about *moral hazard* due to asymmetric information, insinuating that the true IT security level of a policyholder is obscured from the insurer ex-ante or that security measures are purposely neglected ex-post, once an insurance policy is in place. We challenge these arguments and do not believe moral hazard to be a dominant problem in practice, mainly due to the partial non-insurability of cyber and usual policy exclusions in case of wilfully omitted information, see our discussion in Online Appendix A.1 of [2].

Nevertheless, many academic works have been dedicated to the general problem of combining risk mitigation by investing in prevention measures and risk transfer by purchasing insurance, see Section "Related literature" of [2] and the references therein, in particular [144] for an overview relating to the cyber context. These studies usually centralize on a prospective insurance buyer (or a network of such buyers for studies of interdependent security) and treat insurance as an object whose existence enables the considered agents to add a term relating to (pure) risk transfer to their optimization problem.

However, we observe that on the cyber insurance market a different situation presents itself, namely that (prospective) insurance buyers do not perceive insurance policies as a vehicle for mere risk transfer, but in fact expect risk mitigation services, i.e. *cyber assistance*, as part of these policies (see [104, 107]).

Remark 8 (Cyber assistance)**.** *Several large insurers explicitly emphasize the service-oriented character of their cyber insurance policies and the necessity to include prevention and incident-response services through collaboration with specialized IT security providers (see, e.g., [12]). This view seems to be shared by a majority of (prospective) policyholders, who particularly considers the following two types of risk mitigation services desirable as part of cyber insurance solutions (see [104, 107] and Remark 1 in [2]):*

- ***Pre-incident services*** *which include, e.g., network security, back-up of critical systems and data, anti-malware tools, identity and access management, IT security consulting, employee awareness measures, patch management, and mobile device management.*
- ***Post-incident services*** *which contain, e.g., restoration of data, 24h help hotlines, forensic post-breach services, legal advice, and consulting in case of extortion.*

These two types of services on an individual policy level correspond to reducing the *loss probability* and mitigating the *loss severity* respectively, and therefore quite naturally map to the theoretical concepts of *self-protection* and *self-insurance*, see the references in [2], in particular [40] for an excellent introduction to these concepts.

While these types of service may be more relevant for SMEs without specialized IT departments than for large corporations with a high maturity of IT security systems, we additionally advocate for a third type of service relevant to all policyholders, namely using the insurer's portfolio knowledge to install warning mechanisms about imminent threats to reduce loss probabilities across the portfolio (for details, see [2]).

In summary, compared to previous studies we perceive the insurer's role as more central in this case: By setting the price for both risk transfer and risk mitigation, the policyholder can be steered towards a combination of both activities which is optimal for the insurer, considering the balance between her combined income from risk premium / premium for cyber assistance and her overall (portfolio) risk. This does not only provide interesting insights on how to price such combined policies from an insurer's viewpoint, but also entails relevant implications for insurance buyers (namely which price structure a buyer is offered based on his own characteristics and the rest of the portfolio) and for the overall market (namely whether (or rather how) such policies can help alleviate the aforementioned insurance gap).

Before we detail the mathematical framework we have used to approach these questions, we emphasize that we essentially study an insurer's choice of premium which implicates directing buyers towards a certain level of risk mitigation. This general idea is of course not limited to the cyber insurance domain and we briefly review some price discrimination approaches in the cyber context.

Remark 9 (Price discrimination in cyber insurance).

- ***Bonus-Malus system:*** *[144] propose a* Bonus-Malus system *for cyber insurance to incentivize investment in cyber risk mitigation through adjusting the premium of risk transfer. While a Bonus-Malus system may theoretically yield the desired mechanisms (namely to alleviate moral hazard and incentivize the insured to adopt more self-mitigation measures in addition to purchasing insurance), it bases the risk premium purely on the insured's own loss history, which in our view is not a suitable approach for a risk such as cyber which exhibits very low frequency, but high severity losses. This is also emphasized in, e.g., [92], where it is formally shown in the context of screening mechanisms that* post-screening *is not effective at all in a context where losses are rare.*
- ***Risk preference design:*** *[96] use the framework of* risk preference design *to enable the insurer to design incentive-compatible cyber insurance contracts by reshaping the insurance buyers' risk perception. This theoretically allows to quantify and locally reduce the "intensity of moral hazard" and may steer the insurance buyer(s) towards adopting a level of security measures which is closer (compared to a "full-information benchmark") to the ideal action from the insurer's viewpoint. While this is an interesting mechanism for academic studies, its complexity (as it is, e.g., based on the manipulation of a probabilistic distribution of risk preference types in the population) most likely prohibits a real-world application.*
- ***Credibility theory:*** *In our view, a promising approach for price discrimination is the*

framework of credibility theory[26]*: The most common application of combining statistical information about the individual risk and the whole collective is in our view particularly helpful in the cyber context, where data on individual losses is rare, but at the same time risks cannot be assumed homogeneous (see Remark 5). In general, the idea of combining several prior sources of information with varying degree of* credibility *promises broader applicability in the cyber context, e.g. in order to calibrate models by combining statistical information from scarce available data with expert estimations or incorporating data of different maturity (assuming that due to high non-stationarity, older data may be more plentifully available, but less representative), see also Section 5.*

4.2 Measuring and comparing risks

We recall from Section 2.1 that the underlying "object" of an insurance policy is an insured risk, represented by a non-negative random variable X with corresponding c.d.f. F_X. Therefore, in order to study the pricing of policies combining risk transfer and risk mitigation as outlined above, we need a mathematical framework which incorporates formalized answers to the two questions:

1. How can a risk X be *measured*, i.e. mapped meaningfully to some corresponding numerical value (to be interpreted e.g. as monetary units)?
2. How can two risks X and Y be meaningfully *compared*, i.e. when can one risk be classified as "more / less risky" than another?

The first question is fundamental in an insurance context, as it relates to determining the amount of regulatory capital an insurance company needs to hold for a risk (or a collective of risks) it has underwritten. Due to corresponding regulatory frameworks, in practice the most commonly used risk measures are

- *Value-at-Risk:* $\mathbf{VaR}_{1-\gamma}(X) := \inf\big\{x \in \mathbb{R} : \mathbb{P}(X \leq x) \geq 1-\gamma\big\}$, $\gamma \in (0,1)$, and
- *Average Value-at-Risk:* $\mathbf{AVaR}_{1-\gamma}(X) := \mathbb{E}[X | X \geq \mathbf{VaR}_{1-\gamma}(X)]$, $\gamma \in (0,1)$,[27]

which we also use to illustrate the modelling results for risk management in [1, 3].
A general mathematical approach to risk measurement was introduced in the seminal work [18], who coined the axiomatic definition of a *coherent risk measure* based on four desired properties:

Definition 4 (Coherent risk measure)**.** *A map $\rho : L^1(\Omega, \mathcal{F}, \mathbb{P}) \to \mathbb{R}$ is a coherent risk measure if it has the following properties:*

[26] See [32] for an original seminal reference and e.g. [34] for an extensive introduction to the topic.

[27] Note that the term *Average / Tail / Conditional Value-at-Risk* is often used interchangeably with *Expected Shortfall* $\mathbf{ES}_{1-\gamma}(X) := \frac{1}{\gamma}\int_{1-\gamma}^{1} \mathbf{VaR}_z(X)\mathrm{d}z$, $\gamma \in (0,1)$, although in some cases, e.g. for discrete underlying distributions, there are slight technical differences between them, see [5] for a detailed discussion.

1. Monotonicity: $X \leq Y$ *a.s.* $\implies \rho(X) \leq \rho(Y)$*;*

2. Cash-additivity / translation invariance: $\forall\, m \in \mathbb{R}: \ \rho(X - m) = \rho(X) - m$*;*

3. Convexity: $\forall\, X, Y \in L^1, \forall\, \lambda \in [0,1]: \ \rho(\lambda X + (1-\lambda)Y) \leq \lambda\rho(X) + (1-\lambda)\rho(Y)$*;*

4. Positive homogeneity: $\forall\, \lambda \in [0,\infty): \ \rho(\lambda X) = \lambda\rho(X)$*.*

Note that under property 4., convexity is equivalent to sub-additivity.
A coherent risk measure is called law-invariant ([95]) if it additionally fulfils:

5. Law-invariance: $X \overset{(d)}{=} Y \implies \rho(X) = \rho(Y)$*, where* $\overset{(d)}{=}$ *denotes equality in distribution.*

Many authors have built on and expanded this definition since its conception, see Section "Related literature" of [2] for an overview. In an insurance context, an important sub-class of law-invariant, coherent risk measures, which was first connected to insurance pricing in [133, 134, 135], are *concave distortion risk measures* (DRM). For a non-negative risk X, a DRM can generally be expressed as

$$\rho(X) := \int_0^\infty \psi(\overline{F}_X(x))\mathrm{d}x = \int_0^1 \overline{q}_X(u)\mathrm{d}\psi(u), \tag{6}$$

where $\psi : [0,1] \to [0,1]$ is a distortion function (i.e. non-decreasing with $\psi(0) = 0$ and $\psi(1) = 1$), $\overline{F}_X(x) = \mathbb{P}(X > x)$ is the survival function and $\overline{q}_X(u)$, $u \in (0,1)$, its generalized inverse, called the tail quantile function. This framework is convenient, as it includes a rather explicit expression of the decision maker's risk aversion via the (concave) distortion function, which yields the risk measure as a distorted expectation of X.

Remark 10 (**VaR** and **AVaR** as distortion risk measures)**.** *Both Value-at-Risk and Average Value-at-Risk can be expressed in the form of Equation* (6) *by choosing the distortion functions* $\psi_{\mathbf{VaR}}(u) = \mathbf{1}_{\{u>\gamma\}}$ *(not concave) and* $\psi_{\mathbf{AVaR}}(u) = \min\left\{\frac{u}{\gamma}; 1\right\}$ *(concave), respectively. As concavity of the distortion function corresponds to convexity of the risk measure (see [141]), this corroborates that Value-at-Risk does not fulfil all desiderata of a coherent risk measure, which entails e.g. its well-known undesirable property of potentially penalizing diversification.*

While the framework of concave DRMs thus includes Average Value-at-Risk, we prefer to work with the *proportional hazard transform*, introduced for insurance pricing by [133], represented by the distortion function $\psi(u) = u^r$, $r \in (0,1]$, as it additionally possesses the convenient properties of strict concavity and differentiability everywhere. We provide an overview of (other) popular DRMs and underlying distortion functions in Table 1 of [2].

The second question of comparing risks relates to defining a suitable order relation ("stochastic order") between their distributions to express that one is "more risky" than the other. There exist numerous plausible possibilities to define a (partial) ordering on the space of probability distributions; we recommend the excellent book [103] for an overview and a thorough introduction to the topic. We briefly list the two very common stochastic orders which are used in the context of this work:

Definition 5 ('Usual' and 'stop-loss' stochastic order, see, e.g., [103]).
X is "larger" (more risky) than Y in the 'usual stochastic order', write $X \geq_{st} Y$, if their respective c.d.f.s F_X and F_Y fulfil

$$F_X(t) \leq F_Y(t), \quad \forall t \in \mathbb{R}. \tag{7}$$

This is equivalent to (see [103], Theorem 1.2.8):

$$\mathbb{E}[f(X)] \geq \mathbb{E}[f(Y)] \tag{8}$$

for any non-decreasing function $f : \mathbb{R} \to \mathbb{R}$ for which both sides exist. In the economic literature, a common synonymous expression is that X exceeds Y in 'first-order stochastic dominance', write $X \geq_{FSD} Y$. It is obvious that $\geq_{st}$ is a very strong requirement which is not viable for many applications. A slight relaxation of Equation (8) yields a weaker ordering: X is "larger" (more risky) than Y in 'increasing convex order', write $X \geq_{icx} Y$, if (8) holds for all non-decreasing, convex functions for which both expectations exist. Synonymously, $\geq_{icx}$ is often referred to as 'stop-loss order', write $X \geq_{sl} Y$, as it can be equivalently defined via the following condition for the class of so called 'stop-loss functions' (see [103], Theorem 1.5.7):

$$\mathbb{E}[(X-t)_+] \geq \mathbb{E}[(Y-t)_+] \quad \forall t \in \mathbb{R}, \tag{9}$$

where $(x)_+ = \max\{x, 0\}$. For concrete applications, it is more helpful to characterize $\geq_{icx}$ via an equation involving the c.d.f.s akin to (7). This can be accomplished via the following sufficient criterion, called 'single-crossing condition' (originally due to Karlin–Novikoff [91], see Theorem 1.5.17 in [103]): If $\exists\, t_0 \in \mathbb{R}$ such that

$$\begin{aligned} F_X(t) &\geq F_Y(t), \quad \forall t < t_0, \\ F_X(t) &\leq F_Y(t), \quad \forall t > t_0, \end{aligned}$$

and $\mathbb{E}[X] \geq \mathbb{E}[Y]$, then $X \geq_{icx} Y$.

One may argue that a natural way to compare risks is via the numerical value assigned to them by any of the previously introduced risk measures. Indeed, there are connections between the above stochastic order relations and (concave) distortion risk measures.[28]

Remark 11 (Distortion risk measures and stochastic order, see e.g. [49]).
Let X, Y be non-negative random variables.
Any DRM ρ as in (6) preserves the usual stochastic order, i.e. it holds that

$$Y \leq_{st} X \implies \rho(Y) \leq \rho(X).$$

Any DRM ρ as in (6) with concave distortion function preserves the stop-loss order, i.e.

$$Y \leq_{sl} X \implies \rho(Y) \leq \rho(X).$$

[28] The relations can even be stated in a stronger form than is needed here, namely as equivalencies which allow a characterization of the stochastic orders via ordered risk measures (see [49]). For a more general discussion of such relations, we refer e.g. to [22].

The above notions are very useful to formalize important aspects in our study of insurance policies including risk mitigation (presented in a concise, stylized form here). Assume an insurance buyer is facing a risk X.

- The buyer's risk aversion can be captured via a corresponding DRM $\rho_B(X)$, which immediately implies his willingness-to-pay for an insurance policy (i.e. maximum feasible risk premium). Likewise, the risk aversion of the insurer is analogously captured via a DRM $\rho_I(X)$, which entails her minimum feasible risk premium. Naturally, to make the closure of an insurance contract possible, the insurer needs to be less risk-averse than the buyer.[29]
- The effect of purchasing risk mitigation service through cyber assistance can be represented as inducing a decreasing order in first-order stochastic dominance: The c.d.f. $F_{X,s}$ of X, now additionally equipped with a parameter $s \in [0, \infty)$ referring to the amount of service, can be altered "favourably" (i.e. X can be made "less risky" in the sense of $\leq_{st}$) by increasing the level of cyber assistance. Due to the statements in Remark 11, this is directly reflected in a decrease of the risk measures $\rho_B(X)$ and $\rho_I(X)$.

Example 3 (Self-protection and self-insurance inducing $\leq_{st}$)**.** *In order to illustrate how cyber assistance can induce a stochastic order between risks, consider the following elementary example relating to Remark 8. Let X and Y be two risks taking only two values:*

$$X = \begin{cases} 0 & w.p.\ (1 - p_X), \\ K_X & w.p.\ p_X, \end{cases} \qquad Y = \begin{cases} 0 & w.p.\ (1 - p_Y), \\ K_Y & w.p.\ p_Y, \end{cases}$$

where $K_X, K_Y > 0$ and $p_X, p_Y \in (0, 1]$ are constant. Let X represent the original risk (without cyber assistance) and Y the corresponding mitigated risk (subject to some kind of cyber assistance), respectively. Then:[30]

- *The effect of pre-incident services / self-protection is represented by $K_X = K_Y$ (same loss size) and $p_X \geq p_Y$ (decreased loss probability).*
- *The effect of post-incident services / self-insurance is represented by $p_X = p_Y$ (same loss probability) and $K_X \geq K_Y$ (decreased loss size).*

It is straightforward that in both cases Equation (7) *holds, i.e. $X \geq_{st} Y$.*

This framework for measuring and comparing risks, as recently suggested in a more general context in [23], allows to state loss functions for the insurance buyer and insurer which meaningfully capture the effect of risk mitigation service on the insured risk and the "monetary value" both

[29] Note that the inclusion of a risk-neutral insurer is possible in this framework by choosing the identity as distortion function, i.e. $\psi(u) = u$, yielding $\rho_I(X) = \mathbb{E}[X]$.

[30] We emphasize at this point that this is only a stylized example and in reality, risk mitigation measures do typically not have a strictly disjoint effect on either loss probability or loss size, but rather positive consequences for both; compare the discussion in [2].

parties assign to this risk (and therefore accept as a reasonable premium for risk transfer). In summary, both parties deal with their own bivariate optimization problem:

- The insurer sets prices for combined insurance products by selecting a risk loading $\theta \in [0, \infty)$ (where risk transfer is priced according to the expected value principle, see Example 1) and a price for cyber assistance, represented by a cost share $\beta \in [0, 1]$ of administrative service cost charged to the buyer. The choice $\beta = 1$ thus represents full allocation of risk mitigation cost to the buyer, whereas any choice $\beta < 1$ indicates a subsidy for risk mitigation by the insurer. The insurer thereby balances her risk measure $\rho_I(X)$ of the insured risk against the risk premium and (potentially subsidized) service premium she receives.
- The buyer chooses his level of risk transfer (full or no insurance via a proportional insurance share $\alpha \in \{0, 1\}$) and level of risk mitigation (via a parameter $s \in [0, \infty)$), taking into account his risk measure $\rho_B(X)$ and the price of both activities (which has been determined by the insurer).

Due to its sequential nature – the insurer determines prices for risk transfer and risk mitigation before the insurance buyer decides on his course of action depending on the offered premiums – the interaction between insurer and buyer can be modelled as a Stackelberg game and solved via backward induction (see Section “Interaction between cyber-insurance buyer and insurer” of [2] for details and references).

The obtained results offer interesting interpretations of the results derived in [23] in the cyber insurance context from the buyer’s viewpoint and tackle a novel bivariate problem from the insurer’s viewpoint. This part of the study relates to the types of cyber assistance outlined in Remark 8, namely pre-incident and post-incident services an insurer offers to an individual policyholder. As outlined above, a particularly interesting extension from our viewpoint is the consideration of additional services making explicit use of dependence between policyholders. Before we treat this topic in the next section, we close with a remark on risk assessment.

Remark 12 (Cyber risk assessment as a service)**.** *One may wonder about one feature of cyber insurance we have not yet treated in this section, namely cyber risk assessment. While risk assessment is not usually considered explicitly as a* service *within an insurance policy (contrary to the risk mitigation measures outlined in Remark 8), we have, e.g., emphasized in [4] how an insurer’s risk assessment process can provide beneficial insights about a company’s exposure to cyber threats and thus needs for action with respect to cyber security investment.*

However, to emphasize that studying risk assessment requires a different mathematical framework, we refer back to the elementary Example 3: As outlined therein, risk mitigation services can be represented as affecting the distribution function of the underlying risk (i.e. induce a switch from F_X to F_Y) in a way that makes the risk “smaller” in the sense of some stochastic order (in this case $Y \leq_{st} X$), which can be reflected directly in a loss function via a decrease of the corresponding risk measure ($\rho(Y) \leq \rho(X)$).

The study of risk assessment, on the other hand, relates to the acknowledgement that in reality, neither party has perfect knowledge about the objective "true" distribution F_X of the underlying risk and rather works with a subjective assumption $\hat{F}_X^B$ (buyer) and $\hat{F}_X^I$ (insurer). The cost invested in risk assessment aims at increasing the proximity (using an appropriate measure of distance) of the subjective distributions to (the unknown) F_X. This is much harder to capture within a loss function, compare our extensive discussion and ideas for future research in Online Appendix A.1 of [2].

4.3 Towards systemic cyber resilience

In the last section, we have focused on equipping individual cyber insurance policies with risk mitigation services. This is a relevant topic, as it concerns policies already being offered on the market or increasingly demanded by prospective insurance buyers (see Remark 8 and the references therein), and indicates a potential way to close the persistent market gap and keep cyber risk insurable despite hardening market conditions and the challenging nature of the underlying risk (see previous sections).

We have emphasized repeatedly that for cyber risk modelling, from an insurer's viewpoint it is crucial not to view each policy as a stand-alone object, but instead consider the whole portfolio with its possible dependencies between the underlying risks, e.g. via common vulnerabilities which could expose many companies to a potential cyber incident at the same time. While such dependencies are primarily worrisome, as they hamper diversification and may entail accumulation risk, we would like to highlight the following optimistic viewpoint: In reality, cyber incidents from the same root cause do not necessarily manifest as truly simultaneous losses, but may affect victims at different points in time depending on their exposure, the attack vector, the controls (security measures) they have in place, and not least the detection time. By using their knowledge about observed cyber losses in the portfolio, insurers may be in a prime position to prevent or mitigate the manifestation of further losses by warning their policyholders about imminent threats. We approach the question of how to price such services in [2], namely for some bivariate examples of common dependence mechanisms in the cyber context in the Section "The insurer's problem: portfolio viewpoint" and for an extension to a general multivariate portfolio in Online Appendix A.7.3. Many interesting extensions await for future research.

This is a particularly relevant topic, as it relates to the more general question of how cyber insurers can contribute to the overall *cyber resilience* of businesses and networks. The necessity to shift the focus of discussion from mere cyber security to cyber resilience has received increasing emphasis, e.g. in the recent studies [43, 20, 45] (who all carry the term in their title or among their keywords). As mentioned therein, the need to enhance resilience of cyber systems is also reflected by legal and regulatory bodies, consider, e.g., the European Union's proposal of a "Cyber Resilience Act" ([67]) and the World Economic Forum's white paper on a "cyber resilience index" ([143]). This awareness is driven by the acknowledgement that a complete prevention of cyber incidents is not achievable in today's world, where systems are increasingly interconnected and digital technologies ubiquitous.

There are numerous context-specific definitions of resilience (see, e.g., [87] for an overview). NIST[31] defines cyber resiliency as *"the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources"*, clearly extending beyond the scope of cyber security which is focused on prevention and precaution. On an individual company level, this emphasizes the need for a holistic risk management strategy, in which insurance including both pre- and post-incident services can form one important building block.

On the level of interconnected cyber systems, [20] study different security- and topology-based interventions to improve the resilience of networks and relate them to current regulatory measures regarding, e.g., the protection of critical infrastructure or data protection. In this context of systemic risk, resilience can be related to the system's ability to contain a cyber epidemic (i.e. mitigate the impact of a large cyber event), e.g. regarding the number of affected companies, the total economic loss or the total overall "infection" time of all targets (which could be interpreted e.g. as the duration of a business interruption or until complete system functionality is restored). Similar quantities are used to measure the impact of a large cyber event on an insurance portfolio in [84], who emphasize another interesting point regarding operational capacities: If an insurer offers assistance services in cooperation with specialized IT security providers, their efficiency is limited by the number of experts available at any given point in time. If a large number of policyholders requires assistance within a relatively short time period due to a large cyber event / epidemic, individual loss mitigation and containment of the event may be hampered by the inability to respond to all affected policyholders without delay.

In summary, some academic works already represent the hopeful perspective that insurers may, in the future, take a role akin to "private regulators" for cyber resilience by setting guidelines for cyber risk management via their risk assessment and pricing as well as by offering assistance services aimed at risk mitigation on an individual and systemic level. We close this section with a related remark which indicates that many challenges remain for future research and innovation.

Remark 13 (Insurers as "private cyber regulators" - Theory vs. practice)**.** *While cyber security legislation has evolved significantly over the last decade, the legal uncertainty associated with indefinite terms relating to required technical standards often prohibits legal norms from representing clear-cut, actionable road maps (cf. [20]). Therefore, researchers have examined the role of private companies and – more recently – cyber insurers as "private governors" of cyber security (e.g. [139], Section 2.2 of [20] and the references therein), often concluding that insurers can fill this role by promoting holistic cyber risk management through their risk assessment practices and contractual obligations. This may be particularly true for SMEs who are currently*

[31] The US National Institute of Standard and Technology, see `https://csrc.nist.gov/glossary/term/cyber_resiliency`.

often refused cyber coverage due to lack of sufficient cyber security practices (cf. [45]). Nevertheless, we need to mention that while these ideas are often advocated in an academic context, empirical evidence suggests that insurers are not (yet) capable of fulfilling this role at present. [142] study whether cyber insurers actually fulfil the tasks by which they are often claimed to affect cyber security standards (e.g. conduct security assessments or incentivize investment by premium discounts). They conclude that the available evidence thwarts policymakers' hopes and that "[...] cyber insurance appears to be a weak form of governance at present"*. [102] conduct a large number of interviews with industry professionals and discuss among their main themes the role of cyber insurance in setting minimum security practices. They interpret the market's recent hardening w.r.t. premiums and security standards as corroboration of* "how the insurance industry is [...] somewhat behind the curve in understanding what robust cyber security standards in organisations should be" *and conclude that* "the reported state of technical knowledge at present suggests that asking insurance companies to be the arbiters of good cyber security practice is not appropriate at present"*.*

It is clear that cyber insurers will have to play a central role in companies' holistic cyber risk management and for (systemic) cyber resilience, entailing manifold implications for public governance, critical infrastructures, and in general the functionality of various services at the core of public and private life in today's digitized society. However, as outlined above, many cyber insurance design and implementation questions remain to be tackled – and it seems likely that they will continuously entail new challenges for research and innovation given the dynamic nature of digital technologies and cyber risk. We will outline some principal lines of future research in the final section.

5 Outlook and final remark

Cyber risk and cyber insurance is a highly relevant and dynamic topic, with manifold interesting questions from many disciplines continuously appearing. Relating to the topics tackled in this thesis, we focus on giving an overview of some main research directions in the proximity of actuarial science. We naturally do not claim this list to be exhaustive and refer to the many excellent papers (e.g. [137, 45, 52, 53, 72]) on interdisciplinary aspects of cyber risk research for complementary summaries.

The first task of paramount urgency concerns the availability of reliable and representative data for cyber risk research and cyber insurance practice. On the one hand, this refers to data about individual cyber incidents and corresponding economic consequences which enables a better understanding of the statistical properties of (highly heterogeneous) different types of cyber threats and their impacts. On the other hand and of major importance, this refers to data containing information about dependence between cyber incidents, which allows the calibration of models which go beyond the modelling of marginal distributions under an (implicit) independence assumption. Only the availability of suitable data enables the estimation (and evaluation of robustness) of various proposed cyber (dependence) models and therefore allows insurers to incorporate (dependence) models which are based on statistical evidence into their actuarial workflow. In turn, the ability to base their pricing and risk management on a sound mathematical foundation will enable insurers to better quantify the underlying (accumulation) risk and therefore counteract apprehension about cyber moving towards becoming uninsurable. Finally, as due to the high non-stationarity of the cyber landscape, it is doubtful whether the calibration of models based purely on historical data is feasible at all, the adaption of methods from the area of credibility theory which allow combining statistical evidence from data with input from expert opinions seems to be called for.

Another imminent topic regards the design and pricing of novel cyber insurance products including cyber assistance. While we have first approached the topic of pricing systemic risk mitigation services through bivariate applications in [2], this is only the starting point for many related research questions. From a mathematical viewpoint, this encompasses, e.g., extensions to more multivariate examples, the modelling of a three-party interaction additionally including a re-insurer, or questions about optimal risk sharing. Taking a wider perspective, the design of such insurance products carries many economic and legal challenges regarding collaborations between insurers and service providers and in particular – as outlined above – the development of emergency response workflows which are themselves resilient in the face of a large cyber event. Regarding the resilience of cyber networks, helpful qualitative indications for the effective design of real-world systemic interventions can be drawn from studies of cyber epidemics on interdependent networks. Many interesting extensions of the available studies, as well as approaches to making these results more applicable in practice, remain open for future analysis.

In the cyber context, many actuarial tasks remain challenging and require an interdisciplinary approach. This includes risk assessment, where several aspects of the identification of relevant covariates for a company's cyber risk exposure and their subsequent incorporation into a stochastic actuarial model remain unresolved. Hereby, an important issue from a technical viewpoint is the effective quantification (and monitoring) of the cyber security landscape of a potentially highly complex enterprise for actuarial applications. Another pressing challenge for actuaries is cyber insurance pricing, where future mathematical research imminently needs to address the question of pricing dependent risks, a topic which to the best of our knowledge has not been comprehensively addressed, yet.

The problem of designing optimal cyber insurance contracts in practice is a multidimensional and complex, yet increasingly important, task. Actuarial considerations have to be balanced with customer needs as well as legal and governance perspectives to develop products which increase the welfare of all parties and contribute to the overall cyber resilience of systems, processes and the economy as a whole. Interdisciplinary research is needed to investigate welfare implications of cyber security, cyber resilience, and cyber insurance, with the overarching aim of developing recommendations and guideposts for the corresponding regulatory environment and of determining the role cyber insurers can – and should – occupy within the cyber resilience ecosystem. While the design of optimal cyber insurance solutions for companies indirectly affects each private person by enhancing the resilience of businesses and systems omnipresent in daily life, a related open task concerns the development of cyber insurance products for the private customer segment.

In summary, it is uncontested that many interesting and relevant tasks remain for further investigation and the need for comprehensive innovation in the cyber insurance domain will remain pervasive in the future. We close this introduction with the following remark to emphasize once more that the investigation of actuarial challenges related to cyber risk – which are at the core of the development of novel cyber insurance solutions – is not a topic that only mathematicians and actuaries are interested in for professional reasons, but that everyone should care about.

Previously, we have introduced the fundamental task of insurance as making random adverse outcomes quantifiable as risks and enabling individuals and businesses to transfer or share otherwise prohibitively large risks, thereby in a sense shaping public and private life. It is not far-fetched to state that without the existence of adequate (re-)insurance solutions, at this very moment no commercial airplanes and cargo ships would be travelling the air and seas, no wind turbines would be spinning to enable the transition to green energies, and many current home owners and car drivers would have to refrain from the eponymous activities. It is clear that the ubiquitous dependence on information technology – which has long become the core of most systems, processes, and activities all of us rely on in daily life – along with the manifold benefits and opportunities it carries, exposes economy and society to potentially catastrophic adversities in case of failure or malicious exploitation. This realization renders the often-asked question of *whether* cyber risk is insurable somewhat moot, as a negative answer would essentially entail the need to massively curtail or abandon all insurance and – as a consequence – business activity. Instead, it must lead to the question of *how* the resulting risks can be assessed, managed, and mitigated, and which capacities insurers must develop to be able to leverage their existing unique expertise in order to hereby play the key role that befits them.

References

[1] G. Zeller and M. Scherer. A comprehensive model for cyber risk based on marked point processes and its application to insurance. *European Actuarial Journal*, 12(1):33–85, 2022. https://link.springer.com/article/10.1007/s13385-021-00290-1.

[2] G. Zeller and M. Scherer. Risk mitigation services in cyber insurance: optimal contract design and price structure. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 48(2):502–547, 2023. https://link.springer.com/article/10.1057/s41288-023-00289-7.

[3] G. Zeller and M. Scherer. Is accumulation risk in cyber systematically underestimated? *Working paper submitted for publication, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4353098*, 2023.

[4] G. Zeller and M. Scherer. Cyber Insurance: An integral component of Cyber Risk Management. *FIRM e.V. Yearbook 2021*, pages 22–26 and 124–128, 2021. Available at https://www.firm.fm/papers/.

[5] C. Acerbi and D. Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, 2002.

[6] Advisen and PartnerRe. 2017 Survey of Cyber Insurance Market Trends. *Report, available at https://partnerre.com/wp-content/uploads/2017/10/PartnerRe-2017-Survey-of-Cyber-Insurance-Market-Trends.pdf*, October 2017.

[7] Advisen and PartnerRe. 2018 Survey of Cyber Insurance Market Trends. *Report, available at https://partnerre.com/wp-content/uploads/2018/10/2018-Survey-of-Cyber-Insurance-Market-Trends.pdf*, October 2018.

[8] H. Albrecher and S. Asmussen. Ruin probabilities and aggregrate claims distributions for shot noise Cox processes. *Scandinavian Actuarial Journal*, 2006(2):86–110, 2006.

[9] H. Albrecher and O. Boxma. A ruin model with dependence between claim sizes and claim intervals. *Insurance: Mathematics and Economics*, 35(2):245–254, 2004.

[10] H. Albrecher and J. Kantor. Simulation of ruin probabilities for risk processes of Markovian type. *Monte Carlo Methods and Applications*, 8(2):111–127, 2002.

[11] Allianz Global Corporate & Specialty. Allianz Risk Barometer – Identifying the Major Business Risks for 2020. *Report, available at https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/Allianz-Risk-Barometer-2020.pdf*, January 2020.

[12] Allianz Global Corporate & Specialty. Allianz Risk Barometer 2022. *Report, available at https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/Allianz-Risk-Barometer-2022.pdf*, January 2022.

[13] Allianz Global Corporate & Specialty. Cyber attacks on critical infrastructure. *Expert risk article, available at `https://www.agcs.allianz.com/news-and-insights/expert-risk-articles/cyber-attacks-on-critical-infrastructure.html`*, June 2016.

[14] Allianz Global Corporate & Specialty. Making noise about "silent" cyber. *Global Risk Dialogue Winter/Spring 2019 Edition, `https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/grd/AGCS-GRD-Winter-Spring-2019.pdf`*, March 2019.

[15] Allianz Global Corporate & Specialty. A Guide to Cyber Risk: Managing the Impact of Increasing Interconnectivity. *Report, available at `https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/AGCS-Cyberrisk-report.pdf`*, September 2015.

[16] E. Andersen. On the collective theory of risk in case of contagion between claims. *Bulletin of the Institute of Mathematics and its Applications*, 12(2):275–279, 1957.

[17] J. Andress. *The Basics of Information Security: Understanding the Fundamentals of InfoSec in Theory and Practice.* Syngress, Elsevier Science, 2014.

[18] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

[19] Ausschuss Schadenversicherung der Deutschen Aktuarvereinigung e.V. Ergebnisbericht des Ausschusses Schadenversicherung: Use Case der DAV AG Daten und Methoden zur Bewertung von Cyberrisiken. *Technical Report, available at `https://aktuar.de/unsere-themen/fachgrundsaetze-oeffentlich/DAV-Ergebnisbericht_UseCase%20zur%20Modellierung%20von%20Cyberrisiken.pdf`*, 24.10.2022.

[20] K. Awiszus, Y. Bell, J. Lüttringhaus, G. Svindland, A. Voß, and S. Weber. Building Resilience in Cybersecurity – An Artificial Lab Approach. *arXiv preprint `https://arxiv.org/abs/2211.04762`*, 2022.

[21] K. Awiszus, T. Knispel, I. Penner, G. Svindland, A. Voß, and S. Weber. Modeling and pricing cyber insurance: Idiosyncratic, systematic, and systemic risks. *European Actuarial Journal*, 13:1–53, 2023.

[22] N. Bäuerle and A. Müller. Stochastic orders and risk measures: Consistency and bounds. *Insurance: Mathematics and Economics*, 38(1):132–148, 2006.

[23] S. Bensalem, N. Hernández Santibáñez, and N. Kazi-Tani. Prevention efforts, insurance demand and price incentives under coherent risk measures. *Insurance: Mathematics and Economics*, 93(3):369–386, 2020.

[24] B. Berliner. Large risks and limits of insurability. *The Geneva Papers on Risk and Insurance*, 10(37):313–329, 1985.

[25] Y. Bessy-Roland, A. Boumezoued, and C. Hillairet. Multivariate Hawkes process for cyber insurance. *Annals of Actuarial Science*, 15(1):14–39, 2021.

[26] C. Biener, M. Eling, and J.H. Wirfs. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 40(1):131–158, 2015.

[27] G. Blom. Harald Cramér 1893-1985. *The Annals of Statistics*, 15(4):1335–1350, 1987.

[28] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, 2006.

[29] R. Böhme, S. Laube, and M. Riek. A fundamental approach to cyber risk analysis. *Variance*, 12(2):161–185, 2019.

[30] J. Bolot and M. Lelarge. A new perspective on internet security using insurance. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*, pages 1948–1956, 2008.

[31] A. Bouveret. Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment. *IMF Working Paper Series, WP/18/143, International Monetary Fund*, June 2018.

[32] H. Bühlmann. Experience rating and credibility. *ASTIN Bulletin: The Journal of the IAA*, 4(3):199–207, 1967.

[33] H. Bühlmann. *Mathematical Methods in Risk Theory*, volume 172 of *A Series of Comprehensive Studies in Mathematics*. Springer Science & Business Media, second printing edition, 2005.

[34] H. Bühlmann and A. Gisler. *A Course in Credibility Theory and its Applications*, volume 317. Springer Science & Business Media, 2005.

[35] M. Carfora, F. Martinelli, F. Mercaldo, and A. Orlando. Cyber Risk Management: An Actuarial Point of View. *Journal of Operational Risk*, 14(4):77–103, 2019.

[36] J. Cebula and L. Young. A Taxonomy of Operational Cyber Security Risks. *Technical Note CMU/SEI-2010-TN-028, Software Engineering Institute, Carnegie Mellon University*, 2010.

[37] V. Chavez-Demoulin, P. Embrechts, and M. Hofert. An extreme value approach for modeling operational risk losses depending on covariates. *The Journal of Risk and Insurance*, 83(3):735–776, 2016.

[38] V. Chavez-Demoulin and J. McGill. High-frequency financial data modeling using Hawkes processes. *Journal of Banking & Finance*, 36(12):3415–3426, 2012.

[39] R. Cohen, J. Humphries, S. Veau, and R. Francis. An investigation of cyber loss data and its links to operational risk. *Journal of Operational Risk*, 14(3):1–25, 2019.

[40] C. Courbage, B. Rey, and N. Treich. Prevention and precaution. In *Handbook of Insurance*, volume 21, pages 185–204. Springer, New York, NY, 2013.

[41] CRO Forum. CRO Forum Concept Paper on a proposed categorisation methodology for cyber risk. *Technical Report, available at `https://www.thecroforum.org/wp-content/uploads/2016/06/ZRH-16-09033-P1_CRO_Forum_Cyber-Risk_web.pdf`*, June 2016.

[42] Cybersecurity Ventures, sponsored by eSentire. 2022 Official Cybercrime Report. *Technical Report, available at `https://www.esentire.com/resources/library/2022-official-cybercrime-report`*, 2022.

[43] M. Dacorogna, N. Debbabi, and M. Kratz. Building up cyber resilience by better grasping cyber risk via a new algorithm for modelling heavy-tailed data. *European Journal of Operational Research*, 2023.

[44] M. Dacorogna and M. Kratz. Special Issue “Cyber Risk and Security”. *Risks*, 10(112), 2022.

[45] M. Dacorogna and M. Kratz. Managing cyber risk, a science in the making. *Scandinavian Actuarial Journal*, 2023.

[46] D.J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods.* Springer New York, second edition, 2003.

[47] A. Dassios and H. Zhao. Ruin by dynamic contagion claims. *Insurance: Mathematics and Economics*, 51(1):93–106, 2012.

[48] N. Debbabi, M. Kratz, and M. Mboup. A self-calibrating method for heavy tailed data modelling. Application in neuroscience and finance. *arXiv preprint `https://arxiv.org/abs/1612.03974v2`*, 2017.

[49] J. Dhaene, S. Vanduffel, M. Goovaerts, R. Kaas, Q. Tang, and D. Vyncke. Risk measures and comonotonicity: a review. *Stochastic Models*, 22(4):573–606, 2006.

[50] B. Edwards, S. Hofmeyr, and S. Forrest. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity*, 2(1):3–14, 2016.

[51] I. Ehrlich and G. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, 1972.

[52] M. Eling. Cyber risk and cyber risk insurance: Status quo and future research. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 43:175–179, 2018.

[53] M. Eling. Cyber risk research in business and actuarial science. *European Actuarial Journal*, 10(2):303–333, 2020.

[54] M. Eling and K. Jung. Copula approaches for modeling cross-sectional dependence of data breach losses. *Insurance: Mathematics and Economics*, 82:167–180, 2018.

[55] M. Eling and K. Jung. Heterogeneity in cyber loss severity and its impact on cyber risk measurement. *Risk Management*, 24(4):273–297, 2022.

[56] M. Eling, K. Jung, and J. Shim. Unraveling heterogeneity in cyber risks using quantile regressions. *Insurance: Mathematics and Economics*, 104:222–242, 2022.

[57] M. Eling and N. Loperfido. Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: Mathematics and Economics*, 75:126–136, 2017.

[58] M. Eling, M. McShane, and T. Nguyen. Cyber risk management: History and future research directions. *Risk Management and Insurance Review*, 24(1):93–125, 2021.

[59] M. Eling, D. Ning, and R. Ibragimov. Time dynamics of cyber risk. *Conference Paper, American Risk and Insurance Association (ARIA) Annual Meeting 2022*, 2022.

[60] M. Eling and W. Schnell. Extreme cyber risks and the nondiversification trap. *Technical Report / Working Paper*, 2020.

[61] M. Eling and W. Schnell. Ten Key Questions on Cyber Risk and Cyber Risk Insurance. *The Geneva Association — The International Association for the Study of Insurance Economics, Report, available at https://www.genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public/cyber-risk-10_key_questions.pdf*, November 2016.

[62] M. Eling and J.H. Wirfs. *Cyber risk: too big to insure? Risk transfer options for a mercurial risk class*, volume 59 of *I.VW HSG Schriftenreihe*. University of St. Gallen, Institute of Insurance Economics (I.VW-HSG), St. Gallen, 2016.

[63] M. Eling and J.H. Wirfs. What are the actual costs of cyber risk events? *European Journal of Operational Research*, 272(3):1109–1119, 2019.

[64] P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*, volume 33 of *Stochastic Modelling and Applied Probability*. Springer Science & Business Media, 2013.

[65] P. Embrechts, K. Mizgier, and X. Chen. Modeling operational risk depending on covariates: An empirical investigation. *Journal of Operational Risk*, 13(3), 2018.

[66] European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural

persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *OJ L 119*, pages 1–88, 04.05.2016.

[67] European Parliament and Council of the European Union. Proposal for a regulation of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020. *2022/0272 (COD)*, 15.09.2022.

[68] European Parliament and Council of the European Union. Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II). *OJ L 335*, pages 1–155, 17.12.2009.

[69] European Parliament and Council of the European Union. Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011. *OJ L 333*, pages 1–79, 27.12.2022.

[70] European Systemic Cyber Group. Systemic cyber risk. *Technical Report, EU catalogue No DT-04-20-113-EN-N*, February 2020.

[71] M. Fahrenwaldt, S. Weber, and K. Weske. Pricing of cyber insurance contracts in a network model. *ASTIN Bulletin*, 48(3):1175–1218, 2018.

[72] G. Falco, M. Eling, D. Jablanski, V. Miller, L. Gordon, S. Wang, J. Schmit, R. Thomas, M. Elvedi, T. Maillart, et al. A research agenda for cyber risk and cyber insurance. In *Workshop on the Economics of Information Security (WEIS)*, 2019.

[73] S. Farkas, O. Lopez, and M. Thomas. Cyber claim analysis using Generalized Pareto regression trees with applications to insurance. *Insurance: Mathematics and Economics*, 98:92–105, 2021.

[74] R. Fisher and L. Tippett. Limiting forms of the frequency distribution of the largest or smallest member of a sample. In *Mathematical proceedings of the Cambridge philosophical society*, volume 24, pages 180–190. Cambridge University Press, 1928.

[75] C. Frey. Cyberrisiken: Neue Aufgaben für Aktuare. In *Digitalisierung, Cyberrisk & KI: Wandel gestalten. Fakten und Meinungen zur DAV/DGVFM-Jahrestagung 2019*, 2019.

[76] D. Fudenberg and J. Tirole. *Game Theory.* MIT press, 1991.

[77] Gesamtverband der Deutschen Versicherungswirtschaft e.V. Im Mittelstand steigt das Interesse an Cyberversicherungen. *GDV Medieninformationen*, 19.07.2022.

[78] Gesamtverband der Deutschen Versicherungswirtschaft e.V. Unverbindlicher Muster-Fragebogen zur Risikoerfassung im Rahmen von Cyber-Versicherungen für kleine und mittelständische Unternehmen. *GDV Musterbedingungen, available at `https://www.gdv.de/gdv/service/musterbedingungen`*, 2019.

[79] B. Gnedenko. Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of mathematics*, 44:423–453, 1943.

[80] G. Grimmett and D. Stirzaker. *Probability and Random Processes.* Oxford University Press, third edition, 2001.

[81] P. Hagen. Industrie gründet eigenen Versicherer. *Süddeutsche Zeitung*, 09.09.2022.

[82] S.J. Hashemi, S. Ahmed, and F. Khan. Probabilistic modeling of business interruption and reputational losses for process facilities. *Process Safety Progress*, 34(4):373–382, 2015.

[83] H. Herath and T. Herath. Copula-based actuarial model for pricing cyber-insurance policies. *Insurance Markets and Companies*, 2(1), 2011.

[84] C. Hillairet and O. Lopez. Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models. *Scandinavian Actuarial Journal*, 2021:671–694, 2021.

[85] C. Hillairet, O. Lopez, L. d'Oultremont, and B. Spoorenberg. Cyber-contagion model with network structure applied to insurance. *Insurance: Mathematics and Economics*, 107:88–101, 2022.

[86] D. Hofmann and S. Wilson. Advancing Accumulation Risk Management in Cyber Insurance: Prerequisites for the development of a sustainable cyber risk insurance market. *The Geneva Association — The International Association for the Study of Insurance Economics, Technical Report, available at `https://www.genevaassociation.org/sites/default/files/report_advancing_accumulation_risk_management_in_cyber_insurance_0.pdf`*, August 2018.

[87] S. Hosseini, K. Barker, and J. Ramirez-Marquez. A review of definitions and measures of system resilience. *Reliability Engineering & System Safety*, 145:47–61, 2016.

[88] R. Ibragimov, D. Jaffee, and J. Walden. Nondiversification traps in catastrophe insurance markets. *The Review of Financial Studies*, 22(3):959–993, 2009.

[89] R. Ibragimov and J. Walden. The limits of diversification when losses may be large. *Journal of Banking & Finance*, 31(8):2551–2569, 2007.

[90] J. Jacobs. Analyzing Ponemon Cost of Data Breach. *Blog Post, available at `https://datadrivensecurity.info/blog/posts/2014/Dec/ponemon/`*, 11.12.2014.

[91] S. Karlin and A. Novikoff. Generalized convex inequalities. *Pacific J. Math.*, 13(4):1251–1279, 1963.

[92] M. Khalili, X. Zhang, and M. Liu. Effective premium discrimination for designing cyber insurance policies with rare losses. In *Decision and Game Theory for Security: 10th International Conference, GameSec 2019, Proceedings*, pages 259–275. Springer, 2019.

[93] F. Knight. *Risk, Uncertainty and Profit*, volume 31 of *Hart, Schaffner and Marx Prize Essays*. Houghton Mifflin, 1921.

[94] M. Kratz. Introduction to Extreme Value Theory: Applications to Risk Analysis and Management. In *MATRIX book series, vol. 2 - 2017 MATRIX Annals - Mathematics of Risk*, pages 591–636. Springer, 2019.

[95] S. Kusuoka. On law invariant coherent risk measures. In *Advances in Mathematical Economics*, volume 3, pages 83–95. Springer, 2001.

[96] S. Liu and Q. Zhu. Mitigating moral hazard in cyber insurance using risk preference design. *arXiv preprint* `https://arxiv.org/abs/2203.12001`, 2022.

[97] F. Lundberg. Approximerad framställning av sannolikehetsfunktionen: Återförsäkering av kollektivrisker. *Almqvist & Wiksell, Uppsala*, 1903.

[98] J. Mai and M. Scherer. *Simulating copulas: Stochastic models, Sampling algorithms, and Applications*, volume 6 of *Series in Quantitative Finance*. World Scientific Publishing, New Jersey and London and Singapore, second edition, 2017.

[99] T. Maillart and D. Sornette. Heavy-tailed distribution of cyber-risks. *The European Physical Journal B*, 75(3):357–364, 2010.

[100] A. Marotta, F. Martinelli, S. Nanni, A. Orlando, and A. Yautsiukhin. Cyber-insurance survey. *Computer Science Review*, 24:35–61, 2017.

[101] B. Meredith-Miller. Report: Top companies saw cyber rate increases over 80% in Q1 2022. *PropertyCasualty360*, 21.06.2022.

[102] G. Mott, S. Turner, J. Nurse, J. MacColl, J. Sullivan, A. Cartwright, and E. Cartwright. Between a rock and a hard(ening) place: Cyber insurance in the ransomware era. *Computers & Security*, 128:103162, 2023.

[103] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*, volume 389 of *Wiley Series in Probability and Statistics*. Wiley, 2002.

[104] Munich Re. Cyber insurance: Risks and trends 2021. *Munich Re Topics Online, available at* `https://www.munichre.com/topics-online/en/digitalisation/cyber/cyber-insurance-risks-and-trends-2021.html`, 11.03.2021.

[105] Munich Re. Cyber insurance: Risks and trends 2020. *Munich Re Topics Online, available at* *https://www.munichre.com/topics-online/en/digitalisation/cyber/cyber-insurance-risks-and-trends-2020.html*, 14.04.2020.

[106] Munich Re. What if a major cyber attack strikes critical infrastructure? *Munich Re Topics Online, available at* *https://www.munichre.com/topics-online/en/digitalisation/cyber/silent-cyber.html*, 22.11.2018.

[107] Munich Re. Cyber insurance: Risks and trends 2023. *Munich Re Topics Online, available at* *https://www.munichre.com/topics-online/en/digitalisation/cyber/cyber-insurance-risks-and-trends-2023.html*, 26.04.2023.

[108] J. Nešlehová, P. Embrechts, and V. Chavez-Demoulin. Infinite mean models and the LDA for operational risk. *Journal of Operational Risk*, 1(1):3–25, 2006.

[109] R. Böhme and G. Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, 2010.

[110] Y. Ogata. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83(401):9–27, 1988.

[111] M. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, 1994.

[112] H. Panjer. Recursive evaluation of a family of compound distributions. *ASTIN Bulletin: The Journal of the IAA*, 12(1):22–26, 1981.

[113] C. Peng, M. Xu, S. Xu, and T. Hu. Modeling and predicting extreme cyber attack rates via marked point processes. *Journal of Applied Statistics*, 44(14):2534–2563, 2017.

[114] C. Peng, M. Xu, S. Xu, and T. Hu. Modeling multivariate cybersecurity risks. *Journal of Applied Statistics*, 45(15):2718–2740, 2018.

[115] G. Quarg. Non-life insurance. *Lecture notes*, April 2019.

[116] M. Raschke. About the return period of a catastrophe. *Natural Hazards and Earth System Sciences*, 22(1):245–263, 2022.

[117] J. Reinhart. Discussion on ‘A comprehensive model for cyber risk based on marked point processes and its applications to insurance’ (Zeller, Scherer). *European Actuarial Journal*, 12(1):87–88, 2022.

[118] S. Resnick. *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling.* Springer Series in Operations Research and Financial Engineering. Springer Science & Business Media, 2007.

[119] T. Rolski, H. Schmidli, V. Schmidt, and J. Teugels. *Stochastic Processes for Insurance and Finance.* Wiley Series in Probability and Statistics. Wiley, Chichester, Reprinted edition, 2001.

[120] S. Romanosky. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity*, 2(2):121–135, 2016.

[121] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones. Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity*, 5(1):1–19, 2019.

[122] H. Schmidli. *Risk Theory.* Springer Actuarial Lecture Notes. Springer, 2017.

[123] C. Schnell. Cyber-Policen: Warum bei den Versicherern Ernüchterung eingekehrt ist. *Handelsblatt*, 23.11.2022.

[124] O. Schonschek. Letzte Rettung. Marktübersicht: Cyberrisiken versichern lassen. *iX*, pages 48–62, 2020.

[125] G. Schwartz and S. Sastry. Cyber-insurance framework for large scale interdependent networks. In *Proceedigns of the 3rd International Conference on High Confidence Networked Systems*, pages 145–154, 2014.

[126] N. Shetty, G. Schwartz, and J. Walrand. Can competitive insurers improve network security? In *Trust and Trustworthy Computing*, volume 6101, pages 308–322. Springer Berlin Heidelberg, 2010.

[127] S. Shetty, M. McShane, L. Zhang, J. Kesan, C. Kamhoua, K. Kwiat, and L. Njilla. Reducing informational disadvantages to improve cyber risk management. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 43(2):224–238, 2018.

[128] J. Silomon. The Düsseldorf Cyber Incident. *Institute for Peace Research and Security Policy at the University of Hamburg, available at https://ifsh.de/en/news-detail/the-duesseldorf-cyber-incident*, 30.09.2020.

[129] A. Sklar. Fonctions de repartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231, 1959.

[130] statista. Number of internet and social media users worldwide as of April 2023. *Available at https://www.statista.com/statistics/617136/digital-population-worldwide/*, April 2023.

[131] W. Turton and J. Robertson. Microsoft Attack Blamed on China Morphs Into Global Crisis. *Bloomberg*, 08.03.2021.

[132] J. Villaseñor-Alva and E. González-Estrada. A bootstrap goodness of fit test for the generalized Pareto distribution. *Computational Statistics & Data Analysis*, 53(11):3835–3841, 2009.

[133] S. Wang. Insurance pricing and increased limits ratemaking by proportional hazards transforms. *Insurance: Mathematics and Economics*, 17(1):43–54, 1995.

[134] S. Wang. Premium calculation by transforming the layer premium density. *ASTIN Bulletin: The Journal of the IAA*, 26(1):71–92, 1996.

[135] S. Wang. Implementation of proportional hazards transforms in ratemaking. *Proceedings of the Casualty Actuarial Society*, 85(1-2):940–979, 1998.

[136] S. Wang, V. Young, and H. Panjer. Axiomatic characterization of insurance prices. *Insurance: Mathematics and Economics*, 21(2):173–183, 1997.

[137] S. Weber and M. Scherer. Challenges in Cyber Risk and Cyber Insurance: Models, Methods and Data. *ENISA - The European Union Agency for Cybersecurity, Research and Innovation Brief, Series: Annual Research and Innovation Briefs, to be published*, 2023.

[138] J. Werner. Risk Aversion. *The New Palgrave Dictionary of Economics, Palgrave Macmillan*, 2008.

[139] S. Wheatley, A. Hofmann, and D. Sornette. Addressing insurance of data breach cyber risks in the catastrophe framework. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 46:53–78, 2021.

[140] S. Wheatley, T. Maillart, and D. Sornette. The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B*, 89:1–12, 2016.

[141] J. Wirch and M. Hardy. A synthesis of risk measures for capital adequacy. *Insurance: Mathematics and Economics*, 25(3):337–347, 1999.

[142] D. Woods and T. Moore. Does insurance have a future in governing cybersecurity? *IEEE Security & Privacy*, 18(1):21–27, 2019.

[143] World Economic Forum in collaboration with Accenture. The Cyber Resilience Index: Advancing Organizational Cyber Resilience. *White Paper*, July 2022.

[144] Q. Xiang, A. Neufeld, G. Peters, I. Nevat, and A. Datta. A Bonus-Malus Framework for Cyber Risk Insurance and Optimal Cybersecurity Provisioning. *arXiv preprint `https://arxiv.org/abs/2102.05568`*, 2022.

[145] M. Xu, G. Da, and S. Xu. Cyber epidemic models with dependences. *Internet Mathematics*, 11:62–92, 2015.

[146] M. Xu and L. Hua. Cybersecurity insurance: Modeling and pricing. *North American Actuarial Journal*, 23(2):220–249, 2019.

[147] Z. Zhan, M. Xu, and S. Xu. Characterizing honeypot-captured cyber attacks: Statistical framework and case study. *IEEE Transactions on Information Forensics and Security*, 8(11):1775–1789, 2013.

[148] Z. Zhan, M. Xu, and S. Xu. Predicting cyber attack rates with extreme values. *IEEE Transactions on Information Forensics and Security*, 10(8):1666–1677, 2015.

A Core Publications

A.1 A comprehensive model for cyber risk based on marked point processes and its application to insurance [1]

Summary

Over the last decade, cyber risk has been a topic of increased interest in academia and practice and the global cyber insurance market has been developed and expanded at ever-increasing speed. As of 2020, cyber incidents were ranked the number one peril to businesses worldwide ([11]) and since, the Covid-19 pandemic and its side effects may have exacerbated the problem ([104]). Insurers are starting to recognize the market opportunity, while still grappling with a firm understanding of this novel type of risk and its underlying drivers. When writing the first draft of this article in 2020, despite a growing body of academic research on cyber risk (modelling), an established approach for insurance applications was still elusive given the various challenges, e.g. scarcity of historical data, non-stationarity, dependence between losses, and difficult impact quantification. The scope of this article therefore comprises four parts:

The first part (Section 2) extensively surveys existing literature at the time, grouping previous work by the main theme of investigation into (early) game-theoretic studies (e.g. [30, 126, 125]), network models and other approaches to modelling interdependence (e.g. [71, 146, 113]), and statistical investigations of the (scarce) available data on cyber incidents (e.g. [50, 57, 63]). While these works provide valuable contributions to the study of cyber risk in general, their applicability for insurers is often limited due to excessive mathematical complexity, prohibitive data requirements, or limited representativeness of the investigated data for real-world cyber insurance portfolios. Furthermore, a summary of the current cyber insurance market at the time is provided, including studies on insurability (e.g. [26]), common coverage features of cyber insurance policies (see, e.g., [121, 7]), and (mostly ad-hoc) pricing approaches in practice ([121]). This part concludes by emphasizing the potential of designing novel cyber insurance solutions that transcend mere risk transfer. As we regarded this topic as particularly interesting, it was chosen as the subject of a subsequent research project (see [2] below).

The second part (Section 3) introduces a holistic understanding of cyber risk based on technical, legal, financial, and actuarial aspects, and scrutinizes cyber risk factors along the classical decomposition of risk into *threat, vulnerability,* and *impact* (e.g. [100]). A classification of cyber incidents along two dimensions is suggested, namely according to the compromise of the three classical information security protection goals *confidentiality, integrity,* and *availability* (e.g. [17]) and according to the *root cause* of the incident. The former distinction is important, as related literature indicates that different kinds of incidents may vary w.r.t. statistical nature ([35]) and economic consequences ([120]). The latter categories allow to include both malicious and non-malicious origins of cyber incidents and to distinguish between idiosyncratic (independent) incidents and systemic events causing dependent incidents. This reflects the fact that –

based on the aspects described above – we understand *common vulnerabilities* (e.g. [29]) as the source of interdependence between cyber incidents. Lastly, this section identifies key characteristics (covariates) which influence a company's exposure and resilience to cyber threats, regarding both the frequency of cyber incidents and their (monetary) impact.

The next part (Section 4) introduces an actuarial model to capture the above-investigated characteristics of cyber risk and to enable an application to cyber insurance pricing and risk management. To summarize, an adequate cyber risk (portfolio) model should encompass different types of cyber incidents and their dynamic evolution, heterogeneity of companies regarding their characteristics and resulting susceptibility to cyber losses, and dependence of cyber incidents via the exploitation of common vulnerabilities. To ensure the applicability of the model to real-world cyber insurance portfolios, the model combines the separate modelling of frequency and severity of cyber losses in a loss distribution approach and is purposely constructed in a modular way. This allows any end-user to easily amend or replace parts of the model, e.g. based on updated future insights on the nature of cyber risk or on the properties of an internal data set, without discarding the overall structure.

We suggest to model the frequency of cyber incidents by using a marked point process approach: On the one hand, idiosyncratic incidents of each type arrive independently at each company according to an inhomogeneous Poisson process with possibly time- and covariate-dependent rate, reflecting the dynamic evolution of cyber threats and the heterogeneity of the portfolio. On the other hand, systemic events arrive according to an inhomogeneous Poisson process, where each arrival carries a (conditionally i.i.d.) mark encoding the subset of the portfolio affected by the event and the (stylized) "strength" of the event. Subsets affected by a common event represent groups of companies with existing common vulnerabilities. To reflect heterogeneity of the portfolio regarding resilience w.r.t. systemic events, the strength of the event is compared to the (stylized) IT security level of each company in the affected subset, leading to an individually thinned arrival process for actual monetary cyber losses vs. cyber incidents. We analyze the properties of the frequency model (e.g. conditional loss probabilities and overdispersion of systemic loss numbers) and interpret them in the cyber context. We suggest to model loss severities by a combination of (time- and covariate-dependent) lognormal and Generalized Pareto distribution based on the findings of [63]. Alternative approaches for certain types of incidents (based on, e.g., [50, 90, 82]) are outlined in the Online Appendix and can be easily substituted due to the modular model structure.

The final part of the paper (Section 5) provides an example of an actuarial application via an extensive simulation study. The parameters are chosen in line with previous findings from the academic literature; an analysis based on suitable empirical data remained as an interesting task for future research (see [19]). On the individual company level, particular emphasis in the analysis is given to the effect of increased IT security on the loss distribution and the resulting insurance premium. On the portfolio level, to emphasize the relevance of including systemic events and resulting accumulation risk, a comparison with the (default) case of purely independent incidents (with identical marginal distributions) is provided. Furthermore, as the involved severity distributions are quite heavy-tailed (based on [63]), the importance of cover limits, as are customary in primary cyber insurance policies, is illustrated numerically.

Erratum

There is a typo in the rightmost part of the last (non-numbered) equation on p. 61. The correct formula for non-negative losses (as considered here) should read

$$AVaR_{1-\alpha}(L) = [...] \overset{(*)}{=} \frac{1}{\alpha}\int_{1-\alpha}^{1} VaR_{\gamma}(L)\mathrm{d}\gamma.$$

This does not have any consequences for the rest of the article, i.e. in all numerical applications and examples the correct formula is used.

Reception

As the topic is of particularly high interest to practitioners, a discussion of the paper by an industry expert was invited and has been published in the same journal issue as this article ([117]). The discussion is generally favourable and emphasizes the extensive literature review and the holistic analysis of cyber risk as apt and useful. The necessity of a model for systemic events is corroborated as *"the most urgent problem of the cyber insurance industry"*. As the main caveat raised is the question about the applicability of the (very general) full extent of the model due to the scarcity of available data, focussing on the modelling of systemic events is suggested.

Since the publication of this paper, the suggested model has been well received by cyber insurance practitioners and researchers. In particular, the German Association of Actuaries (DAV) working group on data and methods for the valuation of cyber risks ("DAV AG Daten und Methoden zur Bewertung von Cyberrisiken") has suggested and implemented the approach introduced in this paper (in a simplified version) as one of two chosen modelling approaches in their recent report on cyber risk modelling ([19]). Therein, a case study on an exemplary insurance portfolio is included, where the parameters of the models are calibrated based on empirical data and previous findings from the literature. The suggested model is furthermore repeatedly mentioned as a suitable approach for cyber risks from common vulnerabilities in the recent comprehensive survey [21] on the modelling and pricing of cyber insurance.

This article has been awarded the *Gauss Prize* for the best paper in the European Actuarial Journal in 2022 by the German Society for Insurance and Financial Mathematics (Deutsche Gesellschaft für Versicherungs- und Finanzmathematik e.V., DGVFM) and the German Association of Actuaries (Deutsche Aktuarvereinigung e.V., DAV). The work contained in this article (prior to publication) has been presented at several scientific conferences and has received the *FiVeG Award 2020 for the Best Junior Presentation* at the *11th CEQURA Conference on Advances in Financial and Insurance Risk Management.*

Individual contributions

I am the main author of this article. The idea of investigating this topic and the choice of model along the stylized statistical characteristics of cyber losses was developed jointly with my supervisor Matthias Scherer, who also made helpful suggestions regarding both content and presentation of the article during our regular discussions. I was responsible for the writing of the manuscript (the whole first draft as well as subsequent drafts based on comments by Matthias Scherer), developing the proofs of all statements contained therein, and the implementation and presentation of the simulation study.

Permission to include the article

Licence to Publish - Open Access

Journal Name:	European Actuarial Journal	(the 'Journal')
Manuscript Number:	EUAJ-D-20-00047R3	
Proposed Title of Article:	A comprehensive model for cyber risk based on marked point processes and its application to insurance	
Author(s) [Please list all named Authors]:	Gabriela Zeller, Matthias Scherer	(the 'Author')
Corresponding Author Name:	Gabriela Zeller	

Licence Applicable to the Article:

Creative Commons licence CC BY: This licence allows readers to copy, distribute and transmit the Article as long as it is attributed back to the author. Readers are permitted to alter, transform or build upon the Article, and to use the Article for commercial purposes. Please read the full licence for further details at - http://creativecommons.org/licenses/by/4.0/

1 Publication

EAJ Association (the 'Licensee') will consider publishing this article, including any supplementary information and graphic elements therein (e.g. illustrations, charts, moving images) (the 'Article'), including granting readers rights to use the Article on an open access basis under the terms of the stated Creative Commons licence.
Headings are for convenience only.

2 Grant of Rights

Subject to editorial acceptance of the Article, it will be published under the Creative Commons licence shown above.

In consideration of the Licensee evaluating the Article for publication, the Author grants the Licensee the non exclusive , irrevocable and sub-licensable right, unlimited in time and territory, to copy-edit, reproduce, publish, distribute, transmit, make available and store the Article, including abstracts thereof, in all forms of media of expression now known or developed in the future, including pre- and reprints, translations, photographic reproductions and extensions.

Furthermore, to enable additional publishing services, such as promotion of the Article, the Author grants the Licensee the right to use the Article (including the use of any graphic elements on a stand-alone basis) in whole or in part in electronic form, such as for display in databases or data networks (e.g. the Internet), or for print or download to stationary or portable devices. This includes interactive and multimedia use as well as posting the Article in full or in part or its abstract on social media, and the right to alter the Article to the extent necessary for such use. Author grants to Licensee the right to re-license Article metadata without restriction (including but not limited to author name, title, abstract, citation, references, keywords and any additional information as determined by Licensee).

If the Article is rejected by the Licensee and not published, all rights under this agreement shall revert to the Author.

3 Copyright

Ownership of copyright in the Article shall vest in the Author. When reproducing the Article or extracts from it, the Author shall acknowledge and reference first publication in the Journal.

4 Self Archiving

Author is permitted to self-archive a preprint and the accepted manuscript version of their Article.

The rights and licensing terms applicable to the version of the Article as published by the Licensee are set out in sections 2 and 3 above. The following applies to versions of the Article preceding publication by the Licensee and/or copyediting and typesetting by the Licensee. Author is permitted to self-archive a preprint and an Author's accepted manuscript version of their Article.

a) A preprint is the version of the Article before peer-review has taken place ("Preprint"). Prior to acceptance for publication, Author retains the right to make a Preprint of their Article available on any of the following: their own personal, self-maintained website; a legally compliant Preprint server such as but not limited to arXiv and bioRxiv. Once the Article has been published, the Author should update the acknowledgement and provide a link to the definitive version on the publisher's website: "This is a preprint of an article published in [insert journal title]. The final authenticated version is available online at: https://doi.org/[insert DOI]"

b) The accepted manuscript version, by industry standard called the "Author's Accepted Manuscript" ("AAM") is the version accepted for publication in a journal following peer review but prior to copyediting and typesetting:

Author retains the right to make an AAM of their Article available on any of the following, provided that they are not made publicly available until after first publication: their own personal, self-maintained website; their employer's internal website; their institutional and/or funder repositories. AAMs may be deposited in such repositories on acceptance, provided that they are not made publicly available until after first publication.

An acknowledgement in the following form should be included, together with a link to the published version on the publisher's website: "This is a post-peer-review, pre-copyedit version of an article published in [insert journal title]. The final authenticated version is available online at: http://dx.doi.org/[insert DOI]".

5 Warranties & Representations

Author warrants and represents that:

a)
  i. the Author is the sole copyright owner or has been authorised by any additional copyright owner(s) to grant the rights defined in clause 2,
  ii. the Article does not infringe any intellectual property rights (including without limitation copyright, database rights or trade mark rights) or other third party rights and no licence from or payments to a third party are required to publish the Article,
  iii. the Article has not been previously published or licensed,
  iv. if the Article contains materials from other sources (e.g. illustrations, tables, text quotations), Author has obtained written permissions to the extent necessary from the copyright holder(s), to license to the Licensee the same rights as set out in clause 2 and has cited any such materials correctly;

b) all of the facts contained in the Article are according to the current body of research true and accurate;

c) nothing in the Article is obscene, defamatory, violates any right of privacy or publicity, infringes any other human, personal or other rights of any person or entity or is otherwise unlawful and that informed consent to publish has been obtained for all research participants;

d) nothing in the Article infringes any duty of confidentiality which Author might owe to anyone else or violates any contract, express or implied, of Author. All of the institutions in which work recorded in the Article was created or carried out have authorised and approved such research and publication; and

e) the signatory who has signed this agreement has full right, power and authority to enter into this agreement on behalf of all of the Authors.

6 Cooperation

a) Author shall cooperate fully with the Licensee in relation to any legal action that might arise from the publication of the Article, and the Author shall give the Licensee access at reasonable times to any relevant accounts, documents and records within the power or control of the Author. Author agrees that the distributing entity is intended to have the benefit of and shall have the right to enforce the terms of this agreement.

b) Author authorises the Licensee to take such steps as it considers necessary at its own expense in the Author's name(s) and on their behalf if the Licensee believes that a third party is infringing or is likely to infringe copyright in the Article including but not limited to initiating legal proceedings.

7 Author List

Changes of authorship, including, but not limited to, changes in the corresponding author or the sequence of authors, are not permitted after acceptance of a manuscript.

8 Corrections

Author agrees that the Licensee may retract the Article or publish a correction or other notice in relation to the Article if the Licensee considers in its reasonable opinion that such actions are appropriate from a legal, editorial or research integrity perspective.

9 Governing Law

This agreement shall be governed by, and shall be construed in accordance with, the laws of the Federal Republic of Germany. The courts of Berlin, Germany shall have the exclusive jurisdiction.

Springer SpringerLink shop

Subjects Services About Us

Permissions

Get permission to reuse Springer Nature content

Springer Nature is partnered with the Copyright Clearance Center to meet our customers' licensing and permissions needs.

Copyright Clearance Center’s RightsLink® service makes it faster and easier to secure permission for the reuse of Springer Nature content to be published, for example, in a journal/magazine, book/textbook, coursepack, thesis/dissertation, annual report, newspaper, training materials, presentation/slide kit, promotional material, etc.

Simply visit SpringerLink and locate the desired content;

- Go to the article or chapter page you wish to reuse content from. (Note: permissions are granted on the article or chapter level, not on the book or journal level). Scroll to the botton of the page, or locate via the side bar, the "Reprints and Permissions" link at the end of the chapter or article.
- Select the way you would like to reuse the content;
- Complete the form with details on your intended reuse. Please be as complete and specific as possible ao as not to delay your permission request;
- Create an account if you haven’t already. A RightsLink account is different than a SpringerLink account, and is necessary to receive a licence regardless of the permission fee. You will receive your licence via the email attached to your RightsLink receipt;
- Accept the terms and conditions and you’re done!

For questions about using the RightsLink service, please contact Customer Support at Copyright Clearance Center via phone +1-855-239-3415 or +1-978-646-2777 or email springernaturesupport@copyright.com.

How to obtain permission to reuse Springer Nature content not available online on SpringerLink

Requests for permission to reuse content (e.g. figure or table, abstract, text excerpts) from Springer Nature publications currently not available online must be submitted in writing. Please be as detailed and specific as possible about what, where, how much, and why you wish to reuse the content.

Your contacts to obtain permission for the reuse of material from:

- books: bookpermissions@springernature.com
- journals: journalpermissions@springernature.com

Author reuse

Please check the Copyright Transfer Statement (CTS) or Licence to Publish (LTP) that you have signed with Springer Nature to find further information about the reuse of your content.

Authors have the right to reuse their article's Version of Record, in whole or in part, in their own thesis. Additionally, they may reproduce and make available their thesis, including Springer Nature content, as required by their awarding academic institution. Authors must properly cite the published article in their thesis according to current citation standards.
Material from: 'AUTHOR, TITLE, JOURNAL TITLE, published [YEAR], [publisher - as it appears on our copyright page]'

If you are any doubt about whether your intended re-use is covered, please contact journalpermissions@springernature.com for confirmation.

Self-Archiving

- Journal authors retain the right to self-archive the final accepted version of their manuscript. Please see our self-archiving policy for full details: https://www.springer.com/gp/open-access/authors-rights/self-archiving-policy/2124

- Book authors please refer to the information on this link: https://www.springer.com/gp/open-access/publication-policies/self-archiving-policy

European Actuarial Journal (2022) 12:33–85
https://doi.org/10.1007/s13385-021-00290-1

ORIGINAL RESEARCH PAPER

A comprehensive model for cyber risk based on marked point processes and its application to insurance

Gabriela Zeller[1] · Matthias Scherer[1]

Received: 12 August 2020 / Revised: 12 February 2021 / Accepted: 16 July 2021 /
Published online: 17 August 2021

Abstract

After scrutinizing technical, legal, financial, and actuarial aspects of cyber risk, a new approach for modelling cyber risk using marked point processes is proposed. Key covariates, required to model frequency and severity of cyber claims, are identified. The presented framework explicitly takes into account incidents from malicious untargeted and targeted attacks as well as accidents and failures. The resulting model is able to include the dynamic nature of cyber risk, while capturing accumulation risk in a realistic way. The model is studied with respect to its statistical properties and applied to the pricing of cyber insurance and risk measurement. The results are illustrated in a simulation study.

Keywords Cyber risk · Cyber insurance · Emerging risks · Marked point processes · Accumulation risk

1 Introduction

Researchers and practitioners from different disciplines have analysed ‘cyber risk’ and ‘cyber insurance’ from their provenience, among them IT system experts, economists, statisticians, actuaries, etc.; a recent survey of the literature on these topics in business and actuarial science is provided in Ref. [1]. Despite the lack of an established agreed-upon framework, all stakeholders share the opinion that cyber risk is on the rise. This is substantiated by continuously changing and expanding cyber threats [2] and an increasing frequency and magnitude of the financial consequences of cyber incidents [3–6]. The potential consequences of cyber incidents

✉ Gabriela Zeller
gabi.zeller@tum.de

Matthias Scherer
scherer@tum.de

[1] Chair of Mathematical Finance, Technical University of Munich, Parkring 11, 85748 Garching-Hochbrück, Germany

have also been prominently covered by the media; examples being [7–9]. This has lead to various cooperations between academia, industry, and government agencies (e.g. *CISA*[1] in the US or *CiSP*[2] in the UK) with the aim of developing defense strategies against cyber crime and enhancing the overall resilience of IT networks. Corporations have moved their perception of cyber security from a merely technical topic to a larger business risk [10], but this awareness does not yet seem to have translated into widespread institutionalisation of cyber risk management [11]. A recently highlighted aspect is the connection between *Cyber Incidents* and *Business Interruption*, which jointly ranked as the top global business risks in a 2019 survey [12].

One strategy to cope with risk is risk transfer, e.g. via insurance contracts. Parallel to the risk, the demand for cyber insurance solutions has been continuously increasing [10] over the last few years. In spite of its growth, however, today's cyber insurance market is still relatively small compared to the value of the assets that could be impaired by a cyber event [2]. Barriers are not a lack of demand for cyber risk transfer, but rather a number of obstacles that complicate the understanding and quantification of the underlying risk, including the lack of solid data on losses, a fast-paced evolution of cyber risk, and the disparity of data protection laws globally [4, 13].

Despite these challenges, especially in the US an existing market is already established; including underwriters, brokers, and organisations specialized on *cyber data analytics* [14]. Concerning the pricing of cyber risk, however, a surprising finding was published by Romanosky et al. [15]: they systematically analysed cyber policies across the US and found that the main themes used for pricing included *looking to competitors* and *estimation/guesswork*. The ad-hoc nature of cyber policy pricing confirms that a unified quantitative understanding of this new type of risk and its underlying drivers is still at its infancy.

The cyber risk model developed in the present work, designed from an actuarial point of view, constitutes a threefold contribution:

1. The model is based on a holistic approach to cyber risk, systematically describing the underlying risk factors while including information-technological, economic, and actuarial viewpoints.
2. The model is able to capture dependencies and accumulation risk in a realistic way by explicitly taking into account idiosyncratic cyber incidents and systemic cyber events.
3. Using the loss distribution approach, the model can easily be applied in an insurance framework. A simulation study illustrating this application is included.

The remainder of this paper is structured as follows: Sect. 2 carefully reviews the existing literature on cyber risk and identifies key findings for an actuary. Section 3 presents a holistic view on cyber risk, including key characteristics and risk factors.

[1] https://www.cisa.gov.

[2] https://www.ncsc.gov.uk/section/keep-up-to-date/cisp.

A new model is developed and analysed in Sect. 4 and illustrated in a simulation study in Sect. 5. Section 6 concludes and reveals opportunities for further research.

2 Background and literature review

2.1 Literature review

Most papers on cyber risk and cyber insurance are restricted to one particular point of view (e.g. IT security, network modelling, actuarial approaches) and the overall picture remains fragmented. In what follows, we group the existing literature according to the main theme of investigation.

2.1.1 Game-theoretic studies

Bohme and Schwartz [16] studied a unifying framework for modelling cyber insurance and classified existing research approaches of cyber insurance market models. Until 2010, many academic papers were motivated by the study of interdependent security and primarily focused on questions of network security and its relation to the existence of an insurance market, often using game-theoretic approaches (e.g. Refs. [17–22]). Other works concentrated on the correlation properties [23] and monoculture effects [24] of cyber risk and the existence of an insurance market under these conditions.

More recently, a very comprehensive overview of various aspects of cyber insurance was given in Ref. [25], including a classification of existing research approaches with interdependent security according to the underlying insurance market model. While the listed approaches differ in their assumptions, the research aims are quite similar. Most studies focus on the existence of a *Nash Equilibrium* for security investments (e.g. Ref. [26]) and the existence or efficiency of an insurance market (e.g. Refs. [17, 19, 20, 22, 27–30]). Slightly different mathematical approaches include the use of *Bayesian network games* to design optimal cyber insurance contracts [31] or to study the effect of network externality on security adoption [32].

Under quite realistic assumptions, the socially optimal level of security investments cannot be attained in these models, as individuals are incentivised to underinvest [25]. Furthermore, given the availability of cyber insurance, individuals are even more reluctant to invest in self-protection and it is thus generally not possible to design insurance as a means to reach socially optimal levels of investment (e.g. Refs. [19, 20, 22, 27–30]). Some studies thus test whether regulatory actions (e.g. fines or rebates, taxes for low self-protection, or risk pooling arrangements) might enable insurance to incentivise self-protection, reaching conflicting conclusions (e.g. Refs. [17, 19, 22, 28, 29, 33, 34]).

2.1.2 Interdependence and network models

Given that an accepted terminology and framework for cyber risk does not yet exist, some authors concentrated on developing taxonomies and frameworks (e.g. Refs.

[35–37]) or on embedding cyber into the better-known context of operational risk management (e.g. Refs. [38, 39]).

One feature of cyber risk that is commonly regarded as particularly problematic is the lack of independence among the risks/claims, a problem that was addressed using copula approaches in Refs. [40, 41], linear correlations in Ref. [24], and a combination of both in Ref. [23]. More recently, Peng et al. [42] studied the multivariate dependence exhibited by real-world cyber attack data using a Copula-GARCH model. The latter works describe cyber attacks of different types or multivariate nature to be the source of dependence. Peng et al. [43] propose modelling and predicting extreme cyberattack rates by using marked point processes and similarly, studying an empirical data set of breach incidents [44] argue that stochastic processes rather than distributions should be used to model and predict hacking breach incident inter-arrival times and breach sizes. Baldwin et al. [45] find strong evidence of contagion in cyber attacks to different components of a firm's information system using self- and mutually-exciting point processes.

Instead of considering underlying attack rates, studies concerned with cyber insurance seek to quantify the expected monetary losses of an insurer's portfolio. To this end, dependencies between losses can also be captured by considering a model of epidemic spreading on the underlying network of firms. Fahrenwaldt et al. [46] use a (Markovian) SIS-process to model the infectuous spread of a cyber vulnerability and subsequently an adapted counting process for the occurrence of attacks. Xu and Hua [47] use Markovian and Non-Markovian processes for epidemic spreading and propose to use a copula approach to capture the dependence among time-to-infection distributions. Xu et al. [48] study a model of cyber epidemics over complex networks, additionally introducing copulas to capture dependencies between cyberattack events.

2.1.3 Data-driven studies

The lack of publicly available, reliable, and sufficiently large data sets for cyber incidents remains one of the obstacles for sound statistical investigations. Among the best-known data sources on data breaches are the continuously updated *"Chronology of Data Breaches"* dataset by the California-based nonprofit corporation *Privacy Rights Clearinghouse (PRC)*[3] and the *"Open security foundation data loss database"*.[4] The former data was e.g. studied by Edwards et al. [49], with the conclusion that the number of records exposed can be modeled by the log-normal law and the daily frequency can be described by a negative binomial distribution. Somewhat surprisingly, the study found neither size nor frequency of data breaches to exhibit a time trend. Eling and Loperfido [50] use multidimensional scaling and goodness-of-fit tests to analyze the distribution of the data breach information. They show that modelling severity using a log-skew-normal distribution seems adequate and find that different types of data breaches need to be modeled as distinct risk categories.

[3] Available for public download from https://privacyrights.org/data-breaches.

[4] Formerly available for public download from http://datalossdb.org.

Eling and Jung [51] study the cross-sectional dependence of the data breach losses and identify a significant asymmetric dependence of monthly cross-industry losses in four categories by breach types as well as cross-breach type losses in five categories by industries. Farkas et al. [52] analyze heterogeneity of the reported cyber claims through the use of regression trees.

The second database was examined in Ref. [53], who focus on the theft of personal information and report a stable power-law tail distribution of personal identity losses per event. Wheatley et al. [54] combined data from both databases to focus exclusively on large breaches and study maximum breach sizes as well as severity distributions. The best fit is obtained by using a doubly truncated Pareto (Power law) distribution with linearly decreasing shape parameter for breach sizes, with sub-linear growth for the maximum log breach size.

Romanosky [55] uses a (commercial) dataset from *Advisen*, a US-based consultant to the insurance industry, with the aim of examining the composition and costs of cyber events. They conclude that firms may lack a strong incentive to increase their investment in data security and privacy protection and the primary motivation may come from the cyber insurance industry through its use of incentive-based premium reductions.

While the aforementioned papers mostly concentrate on data breaches, Eling and Wirfs [56] has a wider focus: they define cyber risk as a subgroup of operational risk and analyze cyber data from a large operational risk database (*SAS OpRisk Global data*), including a global range of cyber incidents that have occurred over an around twenty-year period and considering actual costs instead of number of affected records only. The frequency of losses is found to be most adequately modelled by a Negative Binomial distribution in a static approach, and a Poisson process with covariate-dependent rate in a dynamic approach based on Ref. [57]. For the loss severity, none of the canonical candidates (exponential, Gamma, log-normal, log-logistic, generalized Pareto, Weibull) were found to accurately model the entire loss data. Promising alternatives were a non-parametric transformation kernel estimation and an extreme value approach, where excesses over a threshold were modelled by a generalized Pareto distribution. The study highlighted the importance of distinguishing between *cyber risks of daily life* and *extreme cyber risks*.

2.2 Background on cyber insurance

Marotta et al. [25] provides an excellent summary of the past, present, and future of the cyber insurance market; as seen in 2017. They report an ongoing growth of available coverage, spurred by rising demand for insurance protection against cyber risks, which in turn is often caused by public coverage of severe cyber incidents [14, 25, 58], the introduction of stricter legislation across the globe [2, 25], and firms' own loss experience [14, 58]. In 2015, the global market for cyber insurance was estimated to be worth around $2 billion in premium, with US business accounting for approximately 90%. At the time, fewer than 10% of all companies had purchased cyber insurance, with typical buyers coming from industries holding large volumes of personal data, such as healthcare and retail, or relying on digitalized technology processes, such as manufacturing and telecommunications. A rapid market growth was projected, with total premium reaching $20+ billion by 2025 [4]. As of today, this estimate still seems realistic, with a global market size of around $7 billion in 2020 [59]. However, despite a strong growth and new insurance solutions being developed continuously, in 2017 the cyber insurance market in the US still had not reached the expected size predicted by optimistic forecasts [25]—thus the question of challenges inhibiting the market development arises.

2.2.1 Challenges and insurability of cyber risk

For the European market's supply side, ENISA et al. [13] identified the lack of solid data on losses, the fast pace of technology evolution, and the lack of adequate reinsurance among the key factors. Regarding the demand side, companies' most often mentioned reasons to refrain from purchasing cyber insurance include high prices [10, 11, 14, 58, 60], lack of availability of desired limits and coverage [14, 58, 60], concerns about numerous exclusions and restrictions [10], and lack of understanding about own exposure [61] or about policy offers [11].

A fundamental question is if, and under what circumstances, cyber risk is insurable at all, given its complex nature. ENISA et al. [13] first examined this question and concluded that cyber might well be an insurable type of risk fulfilling almost all of the considered desiderata. A more detailed analysis based on a dataset from an operational risk database was conducted in Ref. [62] and subsequently adressed in Ref. [63]. Their study identified the main problems to be lack of independence of loss occurrence, presence of information asymmetries, and lack of adequate cover limits. However, they remark that some problematic aspects might be alleviated in the future and thus advocate for systematic data collection, e.g. via platforms for data sharing organised by national regulators or international associations.

2.2.2 Cyber insurance policies: coverage and exclusions

Ignoring the academic question "to be (insurable), or not to be," in practice an immature cyber insurance market has developed and an increasing scope of cyber insurance products is available. The majority of coverage is offered as dedicated cyber coverage [11, 14], with customers frequently shifting from endorsement to standalone policies [58]. The most sought-after types of coverage include cyber-related business interruption, data breaches, cyber extortion, and fund transfer fraud/social engineering [14, 58]. Cyber policies typically cover the most common and costly incidents, including human error, mistakes, and negligence, external attacks by cyber criminals, system or business process failures, and malicious or criminal insiders. Rarely, however, attacks against business partners, vendors, or other third parties are included [10]. All policies generally distinguish between first and third party (liability) losses [15]. A systematic qualitative analysis of cyber insurance policies across the US [15] found a surprisingly strong similarity regarding covered losses, where the ten most commonly covered losses included costs of claims expenses (including legal expenses from penalties, defense, and settlement costs), public relations services, costs of notification of affected individuals, business income loss, data or system restoration, forensic investigation costs, and data extortion expenses. Romanosky et al. [15] points out that the top covered costs are *cleanup costs*, i.e. indirect costs in order to comply with laws, manage the firm's reputation, and reduce further expenses following a breach. Other studies found similar results for the covered types of losses (e.g. Refs. [2, 10, 16, 25]).

Regarding exclusions, Romanosky et al. [15] found more variation between policies, where the most common exclusions stemmed from criminal, fraudulent, or dishonest acts, errors or omissions, intentional violation of a law, criminal investigations or proceedings, and payment of fines, penalties, or fees. Furthermore, hard-to-quantify costs like loss of employee productivity or brand damage are often excluded [10].

Lastly, an important issue to mention is *non-affirmative* or *silent* cyber cover, meaning that cover for cyber incidents may exist for example in traditional property and casualty policies, even though this was not the intention of the underwriter [12]. Misconceptions like this might lead to a dangerous perception gap for insureds [11] who suffer from an illusion of protection as well as insurers who might suffer from (unintentionally written) exposure to cyber risk.

2.2.3 Cyber insurance: risk assessment and pricing in practice

In the US, carriers typically assess an applicant's cyber risk through questionnaires, most of which emphasize the amount and type of data handled by the investigated company, whereas the technical infrastructure and IT security management receive less attention [15]. The sample questionnaire for risk assessment for cyber insurance by the German Insurance Association [64] differentiates between three risk categories primarily according to the annual turnover of a company and, secondarily, according to certain risky business units (e.g. e-commerce or handling of sensitive data), where the number of questions for a candidate increases with increasing risk category.

Regarding pricing, there seem to be large differences between carriers, while surprisingly, some of the recurring themes are reliance on external sources, estimation, comparison with competitors, using underwriter's experience, and adaptation of prices from other insurance lines [15]. Similarly, respondents in Refs. [14, 58] stated that competition between carriers seemed to prevail over actuarial assessment of the cost of risk. Most examined policies in Ref. [15] multiply a base premium by variables relating to standard insurance factors and industry-related factors, where high hazard weightings are assigned to businesses that collect and store a high volume of sensitive data or operate in industries like retail, healthcare, and the financial industry. Finally, premium multipliers are commonly assigned according to the outcome of the questionnaire regarding IT security (e.g. privacy controls, network security controls, existence of an incident response plan). In conclusion, the impression manifests that while insurers are trying to get a better understanding of cyber risk and its drivers, due to the lack of ample reliable data to describe the problem with sufficient statistical precision, as of today pricing often happens on an ad-hoc basis and established quantitative models do not exist, yet.

2.2.4 The potential of cyber insurance: insurance as a service

While traditionally insurance is a means of risk transfer, cyber insurance can potentially offer more than compensation for monetary losses. Many insurers already advertise the services their cyber insurance policies include, e.g. prevention and incident response services or crisis communication support [65]. Moreover, ENISA et al. [13] highlights possible benefits of the development of a cyber insurance market such as the potential to incentivise firms to increase IT security through premium discrimination or the development of a market for security consulting firms that investigate security practices as part of the underwriting process.

Another future topic for insurers concerns arrangements and standards that facilitate sharing data and information about cyber incidents. In order to help corporations to overcome their resentments about sharing such data, it is the insurers' task to demonstrate that pooling data enables them to improve their range of services and design adequate new and transparent products that meet companies' needs [11].

Thus, despite most academic works concluding that in their theoretical frameworks cyber insurance cannot improve social welfare or network resilience, in practice the development of adequate, transparent cyber insurance products and services might entail a number of benefits transcending a mere possibility for companies' cyber risk transfer. In summary, during the last few years research on cyber risk has considerably increased and various aspects have been considered (disjointly). Our work focuses on the viewpoint of actuarial science, but we aim at providing a holistic modelling approach, taking into account both IT security and economic factors.

3 Cyber risk: a holistic view

Cyber risk as a multi-faceted and young risk still lacks an established definition in the (insurance) literature. We therefore introduce key characteristics and risk factors a cyber risk model should comprise.

3.1 Definition and key characteristics

Eling et al. [66] summarizes the origins, consequences, and key characteristics of cyber risk as follows:

> *"Any risk emerging from the use of information and communication technology (ICT) that compromises the confidentiality, availability, or integrity of data or services. [...] Cyber risk is either caused naturally or is man-made, where the latter can emerge from human failure, cyber criminality (e.g. extortion, fraud), cyberwar, and cyber terrorism. It is characterised by interdependencies, potential extreme events, high uncertainty with respect to data and modelling approach, and risk of change."*

Interpreted from the actuarial perspective, the traditional approach of quantifying risk by frequency and severity of incidents, and combining them (potentially using an appropriate dependence structure) to obtain an aggregated loss distribution is complicated for cyber risk. We follow [25, 63] in summarizing the central properties of cyber risk:

- **Absence of historical data:** The novelty of this risk and the absence of an established terminology for cyber incidents makes it difficult for insurers to create a reliable database with information on losses. This is exacerbated by a reporting bias, i.e. companies are often reluctant to reveal incidents in order to avoid reputation damages.
- **Dynamic risk type:** Cyber risk is as non-stationary as the underlying technology and legal framework, which makes the usability of past data for modelling future losses difficult. Among the main features that underscore the dynamic nature of cyber risk are the growing speed and scope of digital transformation, widening sources of vulnerability from hyperconnectivity, and the evolution of threat actors [2].
- **Strategic threat actors:** Cyber losses do not occur in a completely random fashion, as they are often caused by malicious actors with strategic (economic) motives and attack patterns. In 2018, Lewis [6] even described the trend of *cybercrime as a service (CaaS)* encompassing a large diversity and volume of cybercrime offerings, including products (e.g. exploit kits, custom malware) and services (e.g. botnet rentals). Around this, a thriving cybercrime economy has emerged from the related communities, offering for instance product development and technical support.
- **Interdependence/Accumulation risk:** The interconnectedness of IT-systems and the often systemic nature of vulnerabilities induce a dependence structure within and across company networks and the potential for loss accumulation.
- **Interdependence of security:** Another result of the network interdependence are negative externalities regarding security, which within a game-theoretical context might lead to an equilibrium in which all companies underinvest in security and, therefore, the overall network is not sufficiently protected.
- **Difficult impact determination:** Due to the intangible nature of information assets, it is often difficult to quantify the economic consequences of a cyber incident.
- **Information asymmetry:** Cyber insurance exhibits two sorts of information asymmetry: Adverse selection and moral hazard. The former refers to the challenge for an insurer to reliably determine a company's risk exposure, the latter refers to the difficulty of ensuring the risk exposure to be maintained throughout the entire contract period.

As we focus on actuarial questions, we refrain from considering in more detail technological aspects of information security, the economics of cyber security and cybercrime, or the legal framework.

However, one important aspect to be mentioned concerns the role of governments and legislation. For example, in the European Union, the *General Data Protection*

Regulation (GDPR) came into force on May 25, 2018 with fines up to 20 € million or 4% of annual global revenues, emphasizing that the respective legal framework must be considered when modelling the size of cyber insurance claims, as penalties and fines may be included in the coverage. Furthermore, besides setting the legal framework, Anchen [2] argued that the government could help to promote cyber resilience by reshaping incentives and increasing awareness of cyber threats.

3.2 Cyber risk factors

So far, the term *risk* was used informally. Going further, we disintegrate risk into a combination of *threat*, *vulnerability*, and *impact* (c.f. Ref. [25]). A threat is the underlying root cause of the risk, which itself does not necessarily manifest as an *incident*, but is only harmful if there is a corresponding *vulnerability* in the target system. If a threat and an existing vulnerability lead to the occurrence of an incident, the *impact* refers to the consequences, which can be tangible (e.g. direct financial consequences) or intangible (e.g. loss of reputation). The process of risk management classically consists of identifying risks by characterising threats, vulnerabilities, and impacts, analysing risks with regards to the probability and impact of an incident and treating the estimated risks by selecting and applying adequate measures. As outlined in Ref. [25], there are four classical ways of dealing with risks: risk reduction, risk transfer, risk avoidance, and risk acceptance. Clearly, cyber insurance is a tool for risk transfer and a potential incentive for risk reduction.

3.2.1 Threats

In order to assign cyber incidents to a few distinct classes, we recall a quite concise definition of cyber risk originally motivated by the study of operational risk management, namely *"operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems."* [38].

We follow [37] in applying this definition to classify cyber incidents according to three classical information security protection goals: *confidentiality, integrity, and availability* of information assets [67]. Table 1 gives an overview of their definitions and the incident types that compromise each goal.

Of course, these categories are not mutually exclusive; an example combining features of fraud and business interruption is a *Ransomware* attack, i.e. extortion for temporarily withheld data. We nevertheless implement the above distinction, as it is known from data breaches that incidents of different kinds typically show a different statistical nature [68] and, moreover, the economic consequences vary across the incident categories [55]. An incident falling into more than one class could e.g. be assigned partially to both of them according to the losses it entails, e.g. for the Ransomware case, losses from the interruption of operations under BI and losses from ransom payments under FR. Furthermore, we can understand FR as a general class of incidents that cannot be distinctly classified as DB or BI.

Table 1 Classification of three main types of cyber incidents according to the basic information security protection goal they compromise

Type of cyber incident	Abbreviation	Information security goal compromised by this incident type	Definition of information security goal [67]
Data breach or data loss	DB	Confidentiality	Prevention of *unauthorized disclosure* of information
Business interruption	BI	Availability	Prevention of *unauthorized withholding* of information or service
Fraud (or general incident)	FR	Integrity	Prevention/detection of *unauthorized modification or deletion* of information

Until a few years ago, data breaches have been the most observed type of incident, see Sect. 2.1. Recently, however, the potential impact from cyber-related BI has become a major concern [4] whose financial consequences could equal or surpass losses from a data breach. Vice versa, cyber incidents have become the most feared BI trigger [12]. Our classification is quite similar to the definition of a *cyber attack* used in Ref. [13] which, however, only focuses on malicious cyber attacks. In our view, to capture the whole scope of cyber incidents, a second distinction along another dimension, namely the *root cause*, should be made.[5] Here, the following should be distinguished:

- **Targeted attacks**: Malicious attacks that target one firm specifically due to its characteristics and assets. Usually, the attack vector is tailor-made to circumvent the company's defense strategies.
- **Individual failures**: Non-malicious incidents at single firms that happen due to internal or external machine or system malfunction or human error.
- **Untargeted attacks**: Malicious attacks (from an external source) that do not target one firm specifically because of its characteristics, but are opportunistic in the sense that they attack many available targets—usually simultaneously.
- **Mass failures**: Non-malicious events that affect multiple entities simultaneously, such as the failure of a cloud service provider.

Combining incident types and root causes yields the partition of cyber incidents as shown in Table 2.[6] Note that we use terminology that is common in the natural catastrophe context and is applied in the cyber context in Ref. [70]: an incident refers to a single loss, whereas an event can cause many related incidents.[7]

One can further scrutinize motives of individuals or groups for targeting companies via cyber attacks. CRO Forum [69] defines five types of *threat actors* (with corresponding motivation): nation states (strategic), organised criminals (economic), hackers (reputational), hacktivists (political), and insiders. The last group includes *unintentional* insiders, emphasizing that, although malicious attacks are more publicly present, a large share of cyber incidents stems from human error or technical

[5] To avoid confusion, the well-known classification of cyber risk by CRO Forum [69], which distinguishes four types of cyber incidents and four potential root causes, should be mentioned. In this work, we consider their classification's *root causes* in the context of vulnerabilities and denote as *root cause* the actual origin of the incident.

[6] This categorisation also comprises classifications by other sources, e.g. Refs. [51, 55], the *PRC* database, the database of *Advisen* (https://www.advisenltd.com/data/cyber-loss-data/) and the four incident types of Ref. [69].

[7] Note that the common IT terminology of *systemic* vulnerabilities introduced in Sect. 3.2.2 is transferred directly to the terminology of *systemic events* used throughout this work. As this might be reminiscent of the term *systemic risk* used in the finance literature, let us already emphasize that we understand the risk from systemic events in the cyber context as neither the risk of a cascading failure of a whole industry nor a mixture of underlying, non-diversifiable market factors. Rather, we understand that systemic vulnerabilities create common entry points for external threats to the system and therefore introduce the potential for common external shocks to the whole portfolio or parts of it and thus multiple dependent, simultaneous loss occurrences.

Table 2 Examples of cyber incidents/events according to the classification by incident type and root cause

	Idiosyncratic incidents		Systemic events	
	Targeted attack	Individual failure	Untargeted attack	Mass failure
Data breach (DB)	Targeted data theft	Individual unintended data disclosure	Data theft through widespread malware/phishing	Unintended data disclosure at cloud service provider
Business interruption (BI)	Targeted (D)DoS/ransomware attack	Disruption of IT system or process through accidental malfunction	Widespread ransomware attack	Cloud service outage disrupting business services
Fraud/general (FR)	CEO fraud through targeted (spear-)phishing attack	Accidental compromise of database by employee	Widespread ransomware attack or social engineering fraud	Accidental compromise of data stored at cloud service provider

Note that targeted attacks include the special case of supplier attacks and we again understand the category *FR* generally as any incident that compromises data integrity problems. This applies to cyber-related BI [4, 12], data breaches [5, 12], and general cyber incidents [3, 55, 56].

One peculiar type of targeted attacks not yet explicitly mentioned are so-called *supplier attacks*, where a company is not attacked directly but through attacks on supply-chain partners (with potentially weaker defenses). Although this so far only accounts for a minority of incidents, studies have identified a trend of attackers slowly shifting their attack patterns to exploit supply chain partner environments, particularly for industries with mature cybersecurity standards [3], and thus many companies will increasingly seek to extend insurance cover to their supply chains [4].

3.2.2 Vulnerabilities and controls

Threats only manifest as *successful* incidents if there exists an exploitable *vulnerability* in the target system [36]. We distinguish between *symptomatic* and *systemic* vulnerabilities (c.f. Ref. [36, 71]), where the former only affect single firms (e.g. via custom software), while the latter can affect many firms simultaneously (e.g. via a vulnerability in standard software). Especially the second kind is worrisome, as it exposes many potential targets to the same threat and thus could lead to highly correlated and simultaneous losses [36]. From the viewpoint of a company, a vulnerability can be mitigated by establishing adequate *controls*, both technical (e.g. antivirus software) and non-technical (e.g. awareness campaigns [36]).

Investments in cyber security require strategic decisions and cannot be limited to the prevention of cyber incidents, but must also take into account the discovery, investigation, and containment of an attack and the fast recovery of systems to a working state [3]. Many academic works have studied the problem of finding an optimal security level, balancing the cost of controls against the benefits from reduced losses (see Sect. 2.1). We do not further study this problem here, but rather conclude that a firm's IT security level must be a central parameter for an insurance company's risk assessment (as it already is in practice [15]).

Besides *opportunistic* attacks that stem from the opportunity of exploiting an existing vulnerability, we also consider targeted attacks on a specific victim. Thus, further firm characteristics that incentivise such attacks need to be identified. Considering the list of threat actors in the previous section, the following characteristics arise:

- **Industry sector**: Previous studies indicate that both the number and cost of cyber incidents depend on the industry [3, 5, 51, 52, 54, 55], with regulated industries such as healthcare and financial services suffering most. Wheatley et al. [54] mention that the industrial sector as a risk factor may serve as a proxy to identify relatively homogeneous subgroups of companies with respect to their frequency of interaction with consumers and the total volume of personal data they guard.
- **Data**: It is intuitive that indeed the *amount and sensitivity of data* handled by the company is a factor, as especially actors with economic motives will target companies with a high amount of valuable data in order to maximize their economic

gain. In practice, this is already incorporated into insurance pricing via hazard weightings [15].

- **Company size**: Regarding the size of a company, there are different aspects to be considered: Large, publicly known companies are prime targets for threat actors with reputational motives, whereas SMEs are often worse protected due to budget constraints or their smaller awareness for cyber risks.

Eling and Wirfs [56] considered, among others, the company-specific covariates *industry sector* and *size* and found both of them to be highly significant for the frequency of all kinds of cyber incidents in their model.

3.2.3 Impact

Parallel to increasing occurrence rates, the economic consequences of various cyber incidents have recently become more severe, with BI and information loss having the highest monetary impact [3]. For data breach incidents, the average cost could be up to several million USD [5, 10], where the biggest financial consequence is found to be lost business. Quantifying the consequences of cyber incidents is difficult due to the scarcity of historical data and the various (intangible) types of costs. Nevertheless, earlier studies give some indications of cost drivers.

For data breaches, Ponemon Institute LLC [5] find the average *cost per record* to depend on the root cause (malicious attacks vs. system failures and human error) and the industry sector. The latter could be explained by the fact that the rate of lost customers and business depends on the industry, but also by considering the impact of regulation and litigation on breach cost causing highly regulated industries to suffer larger losses [12]. An effect of the company size on the breach cost was reported in Refs. [53] and [55], who developed a model for the log-cost of a data breach depending on the firm’s revenue (as a proxy for size) and the number of compromised records. This is more comprehensive than the well-known *Jacob’s formula* [72], which simply links the log-cost of a data breach to the (log-)number of compromised records. Another amendment was proposed in Ref. [52], who argue that [72] did not yet take into account the cost of *mega data breaches* observed in future years. Finally, adequate controls can not only decrease the probability of a breach, but also its potential consequences: Improvements in data governance programs, presence of incident response plans, and employee training all result in average cost savings in case of a breach [5, 60].

Concluding, there is evidence that for data breaches the cost of an incident depends on the industry sector, the size of the company, the amount of data affected, potentially the type of attack, and controls in place. The statistical findings and distributions used to model the severity of data breaches found in investigations of available databases have been summarized in Sect. 2.1. Note that these findings for data breaches might not necessarily translate to the other incident types, as different types of cyber incidents (e.g. data breaches and privacy violations) are found to display large median cost differences [55].

It is hard to find information on the economic impact of the other two types of incidents studied here, namely BI and fraud. For the former, some sources from the non-cyber domain are available [73–77]. The only sources including indications of which distributions are useful to model economic loss from BI are [77], who finds that the size of yearly BI insurance claims follows a Pareto distribution with an extremely heavy tail and infinite expected claim size, and [75], who suggests modelling BI loss by a *PERT* distribution.

For fraud one might hope to find information in studies on the cost of cyber crime but, unfortunately, the data is usually either aggregated over all types of cyber incidents from malicious sources or focuses on information loss/theft. Thus, there is very little reliable evidence on the actual cost of cyber fraud. Despite indications that cyber risk is quite different from other types of operational risk [56], one way might be to draw on knowledge about the modelling of operational risk as, e.g. the Basel II framework [78] includes *internal fraud* and *external fraud* as event-type categories.

Another option is to refer to the recent work of [56], who study all kinds of cyber incidents (including data breaches as a subset) using a model where the parameters of the distribution of both frequency and severity of cyber incidents might depend on firm-specific and incident-specific covariates as well as time. They resort to an EVT approach to model the severity of events, using the *generalized Pareto distribution (GPD)* to model excesses over a high threshold (the tail of the distribution) and a series of simple parametric distributions (e.g. exponential, Gamma, log-normal) for the body. The GPD with *shape parameter* ξ and *scale parameter* β is of the form

$$GPD_{\xi,\beta}(x) = \begin{cases} 1 - \left(1 + \frac{\xi}{\beta}x\right)^{-1/\xi}, & \text{if } \xi \neq 0, \\ 1 - \exp\left(-\frac{x}{\beta}\right), & \text{if } \xi = 0, \end{cases}$$

for $x \geq 0$ if $\xi \geq 0$ and $x \in [0, -\beta/\xi]$ if $\xi < 0$. They build on the work of [57] to fit a model where the parameters of the GPD may depend on covariates (including time). To the best of our knowledge, their work is the first to model the actual economic loss and to consider general types of cyber incidents instead of only data breaches, thus we incorporate their approach in our framework.

3.3 Properties of a cyber risk model

Before proposing a model for cyber risk, we shortly summarize the properties/stylized facts it should possess given the findings from this chapter:

- Different types of incidents (DB, BI, and FR/general incidents) should be distinguished.
- The model should include idiosyncratic incidents and systemic events, where both categories can include malicious and non-malicious causes. Systemic events stemming from common vulnerabilities are particularly worrisome as they entail accumulation risk.
- Companies should be viewed as heterogeneous, as their exposure and resilience to cyber threats depends on their characteristics. The most relevant such charac-

teristics are the industry sector, the company size, the data handled by the company, and its IT security level.

- The model should be able to capture the dynamic nature of cyber risk, as occurrence rates as well as impact of cyber incidents may change over time.

4 Actuarial model

Considering cyber risk a combination of threat, vulnerability, and impact, from an actuarial viewpoint it remains to translate this understanding into modelling *frequency* and *severity* of losses within a portfolio. After doing so in Sects. 4.2 and 4.3, in Sect. 4.4 we address actuarial questions, before illustrating the model in a simulation study in Sect. 5. Note that the proposed model is purposefully constructed in a modular way, as some of the assumptions and parameter choices might be updated in the future once suitable data is available. Moreover, a user who wants to incorporate properties of an internal data set can refine individual features of the model (or replace parts) without changing the overall structure.

4.1 Insurance portfolio

Consider K firms, labeled $\{1, \dots, K\}$, constituting the portfolio of an insurance company exposed to losses due to cyber incidents (idiosyncratic or caused by systemic events). This typically refers to losses covered by stand-alone cyber policies, but might in some cases include losses that still fall under traditional policies for some insurers (note that a trend towards the elimination of cyber exposure in traditional business is observed, hopefully leading to a clear-cut distinction in the future). Following the findings of Sect. 3, we assume that for each company included into the insurer's portfolio, information about relevant covariates is collected via a questionnaire and public information. Table 3 gives an overview of the characteristics we identified as relevant, the potential to elicit the necessary information from public data or a firm's voluntary disclosure, and a suggestion for their inclusion in a mathematical model.[8] Thus, for each firm $j \in \{1, \dots, K\}$, the vector of covariates

$$x_j = (x_{j1}, \dots, x_{j5})' = (b_j, s_j, d_j, c_j, nsup_j)', \tag{1}$$

is known, yielding a $K \times 5$ *covariate matrix*

[8] We do not claim this list to be exhaustive but stress that all required information can be objectively collected by an insurer. For example, one could argue that for an insurer with a world-wide portfolio, information on a firm's location (jurisdiction) should be added as it might influence the severity (e.g. via fines to be paid following a data breach) as well as, for targeted attacks, the frequency (as data from some countries might be more valuable and therefore a more frequent target) of cyber losses.

$$\mathbf{X} = \begin{pmatrix} x_1' \\ \vdots \\ x_K' \end{pmatrix} = \begin{pmatrix} x_{11} & \cdots & x_{15} \\ \vdots & \ddots & \vdots \\ x_{K1} & \cdots & x_{K5} \end{pmatrix},$$

where each row corresponds to one firm in the portfolio and each column to one of the above covariates in the order given in Eq. (1). We assume for notational convenience that all firms are ordered by sector, i.e. suppose there are B sectors and $K_{\hat{b}}, \hat{b} \in \{1, \dots, B\}$, such that firms in sector $\hat{b}$ exactly correspond to indices $i \in \mathcal{I}_{\hat{b}} := \{1 + \sum_{\ell=1}^{\hat{b}-1} K_\ell, \dots, K_{\hat{b}} + \sum_{\ell=1}^{\hat{b}-1} K_\ell = \sum_{\ell=1}^{\hat{b}} K_\ell\} \subseteq \{1, \dots, K\}$, where $\sum_{\ell=1}^{0} = 0$. This implies that there are exactly $K_{\hat{b}}$ firms in each sector $\hat{b}$ and $K = \sum_{\hat{b}=1}^{B} K_{\hat{b}}$. Additionally, we denote the ordered values of the fourth column of $\mathbf{X}$ as $(c_{[k]})_{k \in \{1,\dots,K\}}$ and additionally define $c_{[0]} = 0$ and $c_{[K+1]} = 1$.[9] Analogously, for each sector $\hat{b} \in \{1, \dots, B\}$, denote the $K_{\hat{b}}$ ordered values of $(c_i)_{i \in \mathcal{I}_{\hat{b}}}$ as $(c^{\hat{b}}_{[k_{\hat{b}}]})_{k_{\hat{b}} \in \{1,\dots,K_{\hat{b}}\}}$, and additionally set $c^{\hat{b}}_{[0]} = 0$ and $c^{\hat{b}}_{[K_{\hat{b}}+1]} = 1$. Thus, on the whole portfolio

$$0 = c_{[0]} \le c_{[1]} \le \dots \le c_{[K]} \le c_{[K+1]} = 1,$$

and on each sector $\hat{b} \in \{1, \dots, B\}$

$$0 = c^{\hat{b}}_{[0]} \le c^{\hat{b}}_{[1]} \le \dots \le c^{\hat{b}}_{[K_{\hat{b}}]} \le c^{\hat{b}}_{[K_{\hat{b}}+1]} = 1.$$

4.2 Loss frequency

We will use the framework of point processes to model the arrival of cyber incidents. This allows to naturally incorporate time- and covariate-dependence of the incident frequency and to distinguish between different types of incidents. A comprehensive overview on point processes is given in Refs. [79, 80], whose notation we use. In the following, all random variables are defined on a suitable probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the state space, $\mathcal{F}$ a σ-algebra on Ω, and $\mathbb{P}$ a probability measure on $(\Omega, \mathcal{F})$. For our purposes, we focus on simple point processes on the non-negative real line, i.e. processes on the state space $[0, \infty)$ interpreted as time, whose corresponding counting process $(N(t))_{t \ge 0} = \Big(|\{i \in \mathbb{N} : t_i \in [0, t]\}| \Big)_{t \ge 0}$ has unit increments, where $|\cdot|$ denotes the cardinality, i.e. the number of elements, of a set.

We recall Table 2 for a classification of cyber incidents according to their incident type and root cause: *Idiosyncratic incidents* (targeted attacks and individual failures) are discussed in the next section, *systemic events* (untargeted attacks and mass failures) are addressed subsequently.

[9] Ties can be ordered arbitrarily.

4.2.1 Idiosyncratic incidents

Idiosyncratic incidents are assumed to occur at each firm independently (from incidents at other firms as well as between types of incidents at the same firm). For these types of incidents, we assume that any incident is *successful* in the sense that it breaches the firms IT security measures and causes a loss. This is reasonable, given that targeted attacks are usually tailor-made against one company and furthermore, the majority of non-successful incidents (*near misses*) of this type might not be monitored or recognized.[10]

We assume the arrival of such incidents of each type $\{\mathrm{DB}, \mathrm{BI}, \mathrm{FR}\}$ at each firm $j \in \{1, \dots, K\}$ to follow an inhomogeneous Poisson process with time- and covariate-dependent rate

$$\lambda^{\cdot,idio}(x_j, t) = \exp\left(f_\cdot(x_j) + g_\cdot(t)\right), \tag{2}$$

where the super-/subscript $\cdot$ stands for one of the incident types $\cdot \in \{DB, FR, BI\}$, the functions $f_\cdot$ additively map (a relevant subset of) the covariates, i.e. $f_\cdot(x) = \alpha_{\lambda,\cdot} + \sum_k f_{\lambda,\cdot,k}(x_{jk})$ for some constant $\alpha_{\lambda,\cdot}$ and $g_\cdot : [0, T] \to \mathbb{R}$ is a measurable function describing the time dependence. The explicit form of the functions $f_\cdot$ and $g_\cdot$ is of course unknown but can be estimated from a suitable data set.[11] The dependence on covariates and time can differ for the three incident types. As an example, if one assumes the rate of data breaches to depend on the covariates x_{j3} (*data*; for targeted attacks), x_{j4} (*IT security*; for failures), and x_{j5} (*number of suppliers*; for supplier attacks) only, this would yield

$$\begin{aligned}\log\left(\lambda^{DB,idio}(x_j, t)\right) &= f_{\lambda^{DB,idio}}(x_{j3}, x_{j4}, x_{j5}) + g_{\lambda^{DB,idio}}(t)\\ &= \alpha_{\lambda,DB} + \sum_{k=3,4,5} f_{\lambda,DB,k}(x_{jk}) + g_{\lambda^{DB,idio}}(t),\end{aligned} \tag{3}$$

where the functions $f_{\lambda,DB,k}$ map factor levels to constants for the ordinal covariates indexed $k \in \{3, 5\}$ (i.e. are naturally measurable) and $f_{\lambda,DB,4}$ is any measurable function of the numerical covariate x_{j4}.

It is clear that for any interval $[\tau_1, \tau_2] \subseteq [0, \infty)$ (set for now $\tau_1 := 0$ and $\tau_2 =: T$), given the covariate matrix $\mathbf{X}$, the number of idiosyncratic incidents of each type arriving at firm j follows a Poisson distribution:

[10] As we assume these attacks to occur due to the firm's characteristics, one might ask if a firm has to simply take its exposure to these types of threats as given. For the occurrence rate of malicious targeted attacks this might be true, but we assume that the impact of a successful attack can be limited by adequate measures (see Sect. 4.3). Furthermore, putting security measures in place mitigates the occurrence of individual failures and potentially implicitly deters targeted attacks as attackers would have to invest more resources to devise an attack vector.

[11] As this ansatz constitutes a standard *generalized additive model*, techniques for parameter estimation are readily available, see, e.g. Ref. [81]. Using the statistical software R, such models can be fit with the function `gam(...,family=poisson)` from the package `mgcv`.

$$\forall j \in \{1, \dots, K\} : N_j^{\cdot,idio}(T) \sim Poi\big(\Lambda_j^{\cdot,idio}(T)\big),$$

where $\Lambda_j^{\cdot,idio}(T) = \int_0^T \lambda_j^{\cdot,idio}(t)\mathrm{d}t$ is the mean measure of the inhomogeneous Poisson process.

As the processes are assumed independent between firms, it follows by superposition that the number of idiosyncratic incidents on the whole portfolio is also Poisson distributed:

$$\sum_{j=1}^{K} N_j^{\cdot,idio}(T) = N^{\cdot,idio}(T) \sim Poi\big(\Lambda^{\cdot,idio}(T)\big), \text{ where } \Lambda^{\cdot,idio}(T) = \int_0^T \Big(\sum_{j=1}^{K} \lambda_j^{\cdot,idio}(t)\Big)\mathrm{d}t.$$

4.2.2 Systemic events

Systemic events cause incidents at multiple firms at the same time and, if of malicious origin, are typically of an opportunistic nature, i.e. a set of firms is affected not because of their specific features or the economic gain attainable from attacking them, but rather due to the availability of an exploitable attack vector against them. This often stems from a common vulnerability, for example a list of bought e-mail-addresses that allows a threat actor to send ransomware to employees of certain firms. In many cases, a common vulnerability would likely affect firms within one industry sector (e.g. if custom software is vulnerable), but of course the common factor can also be unobservable. In any case, to model incidents from systemic events, an extension of the simple point process framework of the previous section is needed. We use the framework of marked point processes, where the process of locations (arrival timepoints of events), now called the *ground process* $N^g(\cdot)$, is a simple[12] point process $\{t_i\}_{i\in\mathbb{N}}$ on the non-negative real line as above, more specifically a non-homogeneous Poisson process with log-rate[13]

$$\log\big(\lambda^{\cdot,g}(t)\big) = g_{\lambda^{\cdot,g}}(t), \tag{4}$$

where the super-/subscript $\cdot$ indicates the event type $\cdot \in \{DB, FR, BI\}$, and again $g_{\cdot}$ is a measurable function of time. Each arrival of the ground process $\{t_i\}_{i\in\mathbb{N}}$ is then equipped with a mark $(m_i, S_i) \in \mathcal{M} \times \mathcal{S}$ consisting of realisations of components $m_i \in \mathcal{M} := [m_{\min}, m_{\max}] \overset{\text{w.l.o.g.}}{=} [0,1]$ and $S_i \in \mathcal{S} := \mathcal{P}_K$, such that the resulting process is a marked point process $\big\{t_i, (m_i, S_i)'\big\}_{i\in\mathbb{N}}$ on $[0,\infty) \times (\mathcal{M} \times \mathcal{S})$, where $\mathcal{M} \times \mathcal{S}$ is called the mark space (for a rigorous definition, see Definition 1 in Online Appendix A.1).

[12] As remarked in Ref. [79], by suitably redefining the marks, any marked point process can be represented as a marked point process on the same state space with a simple ground process N^g.

[13] Of course, the log-link is superfluous in this case and might even seem a bit artificial. However, we decide to use this formulation in order to keep consistent with the previous section, especially as we will see the results from both sections being treated jointly later on.

Remark 1 (Interpretation of mark components)

1. $m_i \in \mathcal{M} = [0, 1]$ describes the strength of an event, where strength can be understood e.g. as effectiveness to overcome IT security measures.[14] This is useful to include, as in reality a wide range of sophistication of attacks exists and capturing their strength allows to quantify the effectiveness of IT security measures and the sensitivity of the expected loss to their improvement.
2. $S_i \in \mathcal{S} = \mathcal{P}_K$ encodes the subset of the portfolio affected by an event.

The two components of the mark are used jointly to determine which firms suffer a loss from a given event, namely those firms included in the affected subset whose security level is lower than the strength of the event. With the above notation, an event $\left(t_i, (m_i, S_i)'\right)$

- arrives at time t_i,
- reaches exactly the firms $\left\{j \in S_i\right\}$, and
- causes a loss in exactly the firms $\left\{j \in S_i^*\right\} := \left\{j \in S_i, c_j < m_i\right\}$.

To characterize a marked point process completely, it remains to specify the conditional distribution of the marks, given the locations of the Poisson ground process N_g (see Proposition 6 in Online Appendix A.1). This is done in the following assumptions whose rationality will be detailed below:

Assumption 1 (Conditional mark distribution)

(A1) The joint mark distribution is independent of the location $t \in [0, \infty)$ and the marks $\{(m_i, S_i)'\}_{i \in \mathbb{N}}$ are independent and identically distributed (iid.).

(A2) The two mark components $\{m_i\}_{i \in \mathbb{N}}$ and $\{S_i\}_{i \in \mathbb{N}}$ are independent, where the distribution of m_i is given by the cdf F_M (with pdf f_M) and the distribution of S_i is given by a (discrete) pmf f_S.

(A3) m_i follows a continuous Uniform distribution on $\mathcal{M} = [0, 1]$.

(A4) The distribution of S_i is generated by distinguishing between general and sector-specific events. Given the event type, firms in the relevant subset are affected with identical probability and independently from each other. More specifically, assume there are r.v. $Z_{ij} \in \{0, 1\}$ – such that $\{j \in S_i\} \iff Z_{ij} = 1$ – whose distribution depends on independent r.v. $G_i \sim Ber(p_G)$ and B_i following some categorical distribution on $\{1, \dots, B\}$ with probability $\{p_1, \dots, p_B\}$ (G_i determines whether the event is sector-specific ($G_i = 1$) or general ($G_i = 0$); B_i determines the affected sector in the former case). Then let

[14] For example, a simple phishing e-mail that would immediately be classified spam is rather weak, whereas a sophisticated exploit designed to circumvent state-of-the-art security systems is rather strong.

$$
\begin{aligned}
\mathbb{P}(Z_{ij} = 1 \mid G_i = 0) \quad &= \; p_{gen} \quad \text{iid. } \forall j \in \{1, \dots, K\}, \\
\mathbb{P}(Z_{ij} = 1 \mid G_i = 1, B_i = \hat{b}) \quad &= \begin{cases} p_{sec} & \text{iid. } \forall j \in \mathcal{I}_{\hat{b}}, \\ 0 & \text{else.} \end{cases}
\end{aligned}
$$

Of course, p_G, p_{sec}, $p_{gen} \in [0, 1]$ and $p_b \in [0, 1]$, $\forall b \in \{1, \dots, B\}$, s.t. $\sum p_b = 1$. We exclude the cases $p_G = p_{gen} = 0$ and $(1 - p_G) = p_{sec} = 0$, which lead to the uninteresting case $\mathbb{P}(S_i = \emptyset) = 1$.

A concrete distributional assumption for a model should ideally be backed by empirical evidence. As this is currently not possible due to data scarcity in the cyber domain, we stick to the principle of imposing as little (unknown) prior information as possible. This justifies (A1) and (A2), as we do not have any evidence that would suggest deviating from iid., to introduce any particular dependence, neither between locations and marks, nor between the components of the mark. Similarly, regarding (A3), one might intuitively rather assume the number of very weak attacks (such as easily recognizable spam e-mails) to be higher than the number of very sophisticated attacks. However, as we do not have statistical evidence that would allow to choose a particular distribution, we use a Uniform distribution (maximum entropy distribution among all continuous distributions on a bounded interval [82]). Considering (A4), several industry experts have highlighted in conversations with us the importance of industry sector-specific systemic events. Thus, we incorporate this idea in our model, while again leaving the distribution as simplistic as possible (conditionally iid. Bernoulli draws). Furthermore, note that due to the modular structure of the model, each assumption can be altered or replaced individually if suitable data indicates the necessity, without compromising the general model structure.

4.2.3 Properties of the model

In the following, we detail properties of the model and their interpretation in the cyber insurance context. As proofs mostly rely on standard techniques, they are given in Online Appendix A.3.

Proposition 1 (Distribution of number of incidents and losses) *Under* (*A*4)*, the number of incidents per event* $\{|S_i|\}_{i \in \mathbb{N}}$ *follows a Binomial mixture distribution, i.e.* $f_{|S_i||n,p}(k) = Binom(n, p, k)$ *with*

$$
(n, p) = \begin{cases} (K, p_{gen}) \text{ with weight } (1 - p_G), \\ (K_{\hat{b}}, p_{sec}) \text{ with weight } p_G \, p_{\hat{b}}, \quad \hat{b} \in \{1, \dots, B\}. \end{cases} \tag{5}
$$

Similarly, under (*A*3) *and* (*A*4)*, the number of losses per event* $\{|S_i^*|\}_{i \in \mathbb{N}}$ *follows a Binomial mixture distribution, i.e.* $f_{|S_i^*||n,p}(k) = Binom(n, p, k)$ *with*

$$
(n, p) = \begin{cases} (K^*, p_{gen}) \quad \text{with weight } (1 - p_G)\,(c_{[K^*+1]} - c_{[K^*]}), \quad K^* \in \{0, \dots, K\}, \\ (k_{\hat{b}}^*, p_{sec}) \quad \text{with weight } \; p_G \, p_{\hat{b}} \, (c^{\hat{b}}_{[k_{\hat{b}}^*+1]} - c^{\hat{b}}_{[k_{\hat{b}}^*]}), \;\; k_{\hat{b}}^* \in \{0, \dots, K_{\hat{b}}\}, \; \hat{b} \in \{1, \dots, B\}. \end{cases} \tag{6}
$$

Notice that the distribution of $\{|S_i^*|\}_{i\in\mathbb{N}}$ implies the distribution of $\{|S_i|\}_{i\in\mathbb{N}}$ as the special case where $c_{[k]} = 0, \forall k \in \{1,\ldots,K\}$, i.e. the worst-case scenario where no firm has any IT security measures in place and thus the number of incidents and losses is equivalent.

Proposition 2 (Conditional incident and loss probability) *For a firm $j_1 \in \{1,\ldots,K\}$ in sector b_{j_1}, the probability of being affected by an event, given the information that another firm $j_2 \in \{1,\ldots,K\}$ in sector b_{j_2} has been affected (i.e. the conditional incident probability), is given by*

$$\mathbb{P}(j_1 \in S_i \mid j_2 \in S_i) = \begin{cases} \dfrac{p_{sec}^2\, p_{b_{j_2}}\, p_G + p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} = b_{j_2}, \quad (7a) \\[2ex] \dfrac{p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} \neq b_{j_2}, \quad (7b) \end{cases}$$

where

$$\tilde{p}(b_j) := \mathbb{P}(j \in S_i \mid b_j) = p_G\, p_{b_j}\, p_{sec} + (1-p_G)\, p_{gen} \tag{8}$$

is the (unconditional) incident probability for each firm, given its industry sector.

Likewise, for the conditional loss probabilities,

$$\mathbb{P}(j_1 \in S_i^* \mid j_2 \in S_i^*) = \begin{cases} \dfrac{p_{sec}^2\, p_{b_{j_2}}\, p_G + p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} = b_{j_2},\ c_{j_1} \leq c_{j_2}, \quad (9a) \\[2ex] \dfrac{p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} \neq b_{j_2},\ c_{j_1} \leq c_{j_2}, \quad (9b) \\[2ex] \dfrac{\bar{F}_M(c_{j_1})}{\bar{F}_M(c_{j_2})}\left(\dfrac{p_{sec}^2\, p_{b_{j_2}}\, p_G + p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}\right), & b_{j_1} = b_{j_2},\ c_{j_1} > c_{j_2}, \quad (9c) \\[2ex] \dfrac{\bar{F}_M(c_{j_1})}{\bar{F}_M(c_{j_2})}\left(\dfrac{p_{gen}^2\,(1-p_G)}{\tilde{p}(b_{j_2})}\right), & b_{j_1} \neq b_{j_2},\ c_{j_1} > c_{j_2}, \quad (9d) \end{cases}$$

where the unconditional loss probability is given by

$$\mathbb{P}(j \in S_i^* \mid b_j) = \bar{F}_M(c_j)\,\big(p_G\, p_{b_j}\, p_{sec} + (1-p_G)\, p_{gen}\big) = \bar{F}_M(c_j)\,\tilde{p}(b_j). \tag{10}$$

The above results are interesting from a practical viewpoint: If an insurer is notified about a cyber incident by one of its policyholders (many policies include mandatory immediate notification or even the provision of an immediate-response-team by the insurer), it is worthwhile to find (and warn!) firms with a high conditional probability of having been affected by the same event, thus potentially giving them the chance to avert the actual manifestation in their firm (e.g. by warning employees about a phishing threat or updating vulnerable software). The information about an incident in one firm always has a non-negative effect on the incident probabilities for other firms of the same sector ((7a) vs. (8); a formal proof of this statement is given in Online Appendix A.3), while the effect can go in either direction for firms of different sectors ((7b) vs. (8)). For a detailed illustration, see Fig. 1. The same holds for the information of a suffered loss, i.e. the probability of suffering a loss increases with the knowledge that another firm of the same sector has suffered a loss, and the increase is larger if the harmed firm's IT security level exceeds the one of the firm under consideration.

Analogously to the notation in the previous section, for any interval $[0, T] \subseteq [0, \infty)$, given the arrival process $\left\{t_i, (m_i, S_i)'\right\}_{i \in \mathbb{N}}$ and the covariate matrix $\mathbf{X}$, the number of incidents $\bar{N}_j^{\cdot,syst}$ resp. losses $N_j^{\cdot,syst}$ at each firm follows a Poisson process, where the rate can be obtained by thinning the ground process $N^{\cdot,g}$ of arrivals $\{t_i\}_{i \in \mathbb{N}}$ appropriately (see Ref. [83] and Proposition 5 in Online Appendix A.1). In particular

$$\bar{N}_j^{\cdot,syst}(T) = \sum_{i=1}^{N^{\cdot,g}(T)} \mathbb{1}_{\{j \in S_i\}} \sim Poi\big(\tilde{p}(b_j)\Lambda^{\cdot,g}(T)\big),$$

$$N_j^{\cdot,syst}(T) = \sum_{i=1}^{N^{\cdot,g}(T)} \mathbb{1}_{\{j \in S_i^*\}} \sim Poi\big(\tilde{p}(b_j)\bar{F}_M(c_j)\Lambda^{\cdot,g}(T)\big).$$

Contrary to the previous section, we cannot transition to the portfolio level by simple superposition due to lack of independence between firms. Instead, we express the cumulative number of incidents $\bar{N}^{\cdot,syst}(T)$ resp. losses $N^{\cdot,syst}(T)$ across the entire portfolio for fixed $T > 0$ as a *compound Poisson* distributed r.v.

$$\bar{N}^{\cdot,syst}(T) = \sum_{i=1}^{N^{\cdot,g}(T)} |S_i| \quad \text{and} \quad N^{\cdot,syst}(T) = \sum_{i=1}^{N^{\cdot,g}(T)} |S_i^*|,$$

where $N^{\cdot,g}(T) \sim Poi\Big(\int_0^T \lambda^{\cdot,g}(t)\mathrm{d}t\Big)$ and $\{|S_i|\}_{i \in \mathbb{N}}$ resp. $\{|S_i^*|\}_{i \in \mathbb{N}}$ are iid. mixed Binomial and independent from $N^{\cdot,g}(T)$.[15] Using well-known results for the calculation of the expectation and variance of a compound Poisson r.v. (details in Refs. [84, 85] and Online Appendix A.2), this yields:

Proposition 3 (Overdispersion of systemic incident/loss numbers) *Assume* $K > 1$ *and* $K_{\hat{b}} > 1$ *for at least one* $\hat{b} \in \{1, \dots, B\}$ *with* $p_{\hat{b}} > 0$. *Then, the cumulative number*

[15] The notation $\bar{N}$ and N alludes to the fact that the number of incidents can always be considered a worst-case bound for the number of losses (counterfactual analysis: what had happened if no security was in place at all); in particular, for a given realisation $\{t_i, (m_i, S_i)\}_{i: t_i \in [0,T]}$ always $\bar{N}^{\cdot,syst}(T) \geq N^{\cdot,syst}(T)$.

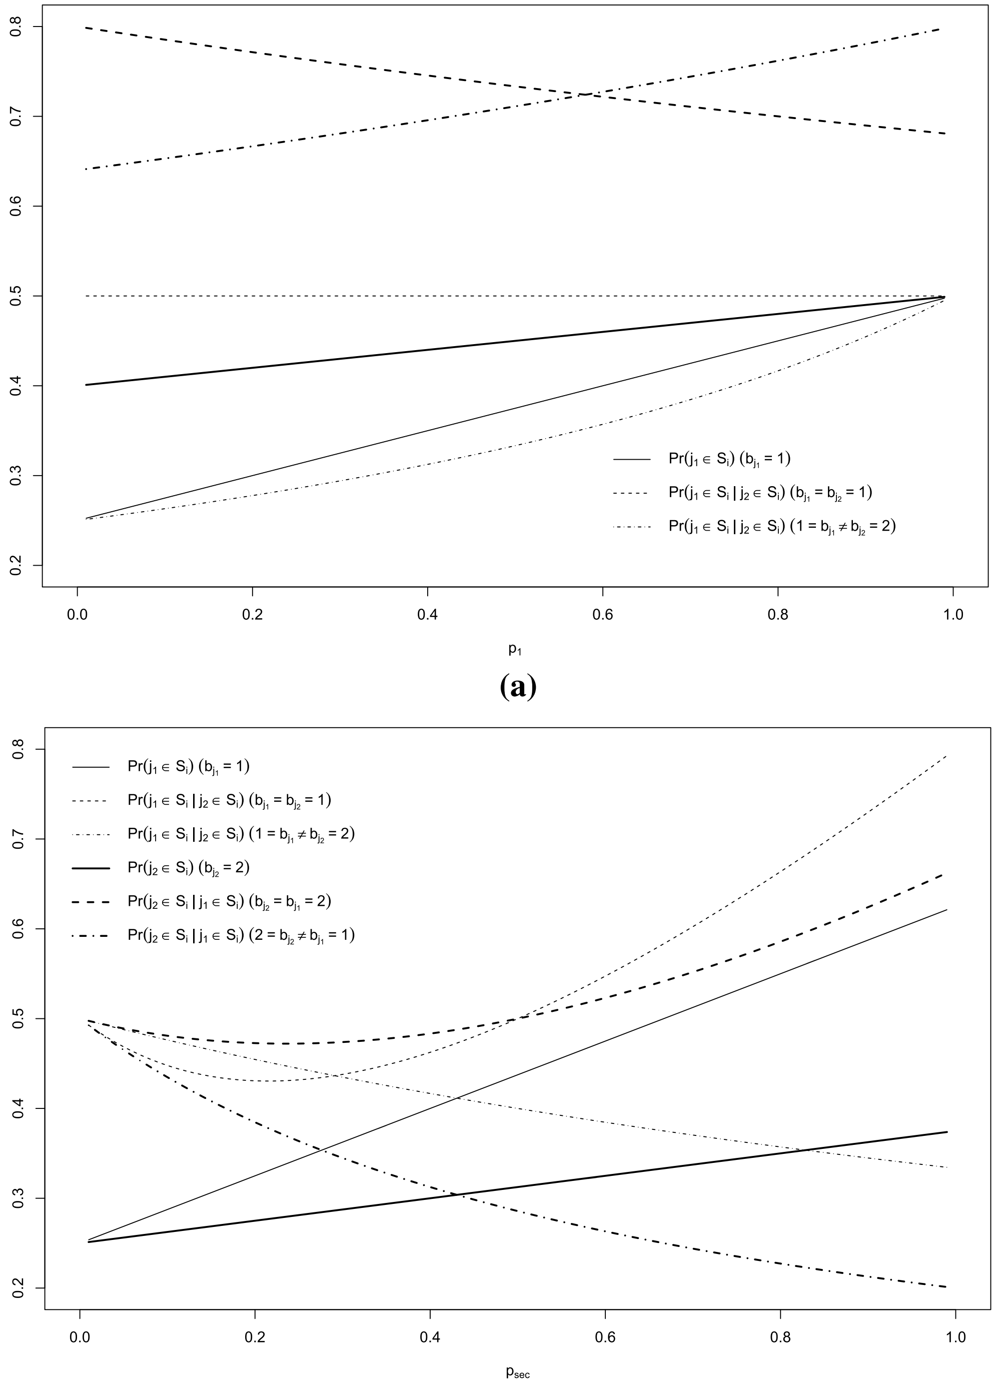

(a)

(b)

Fig. 1 We illustrate the effect of the information that a firm of the same (resp. a different) sector has been affected by an event on the incident probability for just two sectors. (**a**) shows the (un) conditional probabilities for a firm from sector 1 dependent on p_1. The other parameters are chosen as $p_G = p_{sec} = p_{gen} = 0.5$ (thin lines) or $p_G = 0.5, p_{sec} = 0.2, p_{gen} = 0.8$ (thick lines), such that one observes that for $b_{j_1} \neq b_{j_2}$, the conditional probabilities can be above or below the unconditional one (solid line), whereas for $b_{j_1} = b_{j_2}$, the conditioning has a non-negative effect in both cases. Likewise in (**b**), for $p_G = 0.5, p_{gen} = 0.5, p_1 = 1 - p_2 = 0.75$, probabilities for all cases are shown dependent on p_{sec}. Observe again that conditioning on the same sector has a non-negative effect, whereas when conditioning on the other sector, there is an intersection $\mathbb{P}(j_1 \in S_i) = \mathbb{P}(j_1 \in S_i \mid j_2 \in S_i)$ (and $\mathbb{P}(j_2 \in S_i) = \mathbb{P}(j_2 \in S_i \mid j_1 \in S_i)$) at $p_{sec} = 0.4305$

of incidents resp. losses from systemic events is overdispersed, i.e. has a dispersion index (DI = variance-to-mean ratio) exceeding 1.

$$\begin{aligned} DI\big(\bar{N}^{\cdot,syst}(T)\big) &:= \frac{\mathbb{V}ar\big[\bar{N}^{\cdot,syst}(T)\big]}{\mathbb{E}\big[\bar{N}^{\cdot,syst}(T)\big]} = 1 + \frac{(1-p_G)\, p_{gen}^2\, (K^2-K) + p_G\, p_{sec}^2 \sum_{\ell=1}^{B} p_\ell (K_\ell^2 - K_\ell)}{(1-p_G)\, K\, p_{gen} + p_G\, p_{sec} \sum_{\ell=1}^{B} p_\ell\, K_\ell} > 1, \\ DI\big(N^{\cdot,syst}(T)\big) &= 1 \\ &+ \frac{(1-p_G)\sum_{k^*=0}^{K} p_{gen}^2((k^*)^2 - k^*)(c_{[k^*+1]} - c_{[k^*]}) + p_G \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_{sec}^2 p_\ell ((k_\ell^*)^2 - k_\ell^*)(c_{[k_\ell^*+1]}^{\ell} - c_{[k_\ell^*]}^{\ell})}{(1-p_G)p_{gen}\sum_{k^*=0}^{K} k^*(c_{[k^*+1]} - c_{[k^*]}) + p_G p_{sec} \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_\ell k_\ell^* (c_{[k_\ell^*+1]}^{\ell} - c_{[k_\ell^*]}^{\ell})} > 1. \end{aligned}$$

It is well-known that the Poisson distribution exhibits equidispersion ($DI = 1$), while empirical studies in non-life insurance often report overdispersed claim count data and therefore recommend to use alternative distributions (e.g. negative Binomial) to model claim counts. Proposition 3 shows that our construction using a marked Poisson process to allow simultaneous arrivals of several incidents (resp. losses) due to one systemic event likewise is able to introduce overdispersion.

4.2.4 Summary: loss frequency model

Recall that idiosyncratic incidents arrive to each firm $j \in \{1, \dots, K\}$ independently as inhomogeneous Poisson processes with rates $\lambda_j^{\cdot,idio}(t)$, $\cdot \in \{DB, FR, BI\}$. Each incident is assumed to cause a loss. For a fixed time $T > 0$, the overall number of losses caused by idiosyncratic incidents up to this time $N^{\cdot,idio}(T)$ follows a Poisson distribution with rate $\Lambda^{\cdot,idio}(T) = \int_0^T \Big(\sum_{j=1}^{K} \lambda_j^{\cdot,idio}(t) \Big) \mathrm{d}t$.

Systemic events arrive to the portfolio with overall rates $\lambda^{\cdot,g}(t)$. Each arrival t_i carries a mark including the strength of the event m_i and the affected subset S_i. An event at time t_i thus causes incidents in all firms in the set $\{j \in S_i\}$ and causes losses in its subset $\{j \in S_i^*\} = \big\{j \in S_i, c_j < m_i\big\}$. The total number of incidents and losses from systemic events up to time $T > 0$, $\bar{N}^{\cdot,syst}(T)$ resp. $N^{\cdot,syst}(T)$, follow a compound Poisson distribution with mixed Binomial jump sizes.

Aggregating the number of incidents and losses from both root causes on the level of each individual firm translates to aggregating two independent Poisson r.v.

$$N_j^{\cdot}(T) := N_j^{\cdot,idio}(T) + N_j^{\cdot,syst}(T) \sim Poi\big(\Lambda_j^{\cdot,idio}(T) + \tilde{p}(b_j)\bar{F}_M(c_j)\Lambda^{\cdot,g}(T)\big) =: Poi\big(\Lambda_j^{\cdot}(T)\big). \tag{11}$$

On the portfolio level, we aggregate two independent compound Poisson r.v. (one with jumps of constant size 1 and one with mixed Binomial jump sizes), which yields (see Proposition 7 in Online Appendix A.2):

$$N^{\cdot}(T) := N^{\cdot,idio}(T) + N^{\cdot,syst}(T) \overset{d}{=} \sum_{i=1}^{N(T)} Y_i, \tag{12}$$

where $N(T) \sim Poi\big(\Lambda^{\cdot,idio}(T) + \Lambda^{\cdot,g}(T)\big)$ and $\{Y_i\}_{i\in\mathbb{N}}$ are iid., independent of $N(T)$, with mixture distribution

$$F_{Y_i}(n) = \frac{\Lambda^{\cdot,idio}(T)}{\Lambda^{\cdot,idio}(T) + \Lambda^{\cdot,g}(T)}\, \mathbb{1}_{[1,\infty)}(n) + \frac{\Lambda^{\cdot,g}(T)}{\Lambda^{\cdot,idio}(T) + \Lambda^{\cdot,g}(T)}\, F_{|S_i^*|}(n),\ n \in \mathbb{N}_0.$$

Note that so far, we have kept the numbers of losses of different types {DB, FR, BI} separate. In general, assuming independence between them, they could be aggregated into one arrival process, but this may not be desirable as also the loss severity distributions might be different (see Sect. 3.2.3), and therefore for the determination of the portfolio loss their numbers have to be taken into account separately.

4.3 Loss severity

After describing the model of the cumulative number of cyber incidents and losses in the last section, we now turn to their impact, i.e. let $L_{ij} := L_j(t_i)$ be a r.v. describing the non-negative monetary loss caused by a cyber incident reaching firm $j \in \{1, \dots, K\}$ at time $t_i \in [0, T]$. Based on previous findings from academic literature and the arguments in Sect. 3.2.3, we model the body and tail of the loss severity distribution separately and allow the parameters of the distributions to exhibit time- and covariate-dependence. Specifically, for all types of incidents, we suggest using a combination of log-normal and generalized Pareto distribution based on the findings of Ref. [56]. Other promising approaches (e.g. based on Refs. [49] and [72] for DBs or based on Ref. [75] for BIs) are detailed in Online Appendix A.4.

As we do not rely on empirical data, we first need to set a threshold between body and tail of the to-be-constructed distribution. Therefore, we first assume an underlying log-normal distribution $\widetilde{L}^{\cdot}_{ij} \sim LN(\mu^{\cdot}_{ij}, \sigma^{\cdot})$ and select a high quantile as threshold, e.g. set $u^{\cdot}_{ij} = q_z(\widetilde{L}^{\cdot}_{ij})$ with e.g. $z = 0.95$. Given the threshold, construct the density $f_{L^{\cdot}_{ij}}$ of the loss distribution as[16]

$$\begin{aligned}
f_{L^{\cdot}_{ij}}(l) &= \begin{cases} z\, f_{TruncLN}(l;\mu^{\cdot}_{ij}, \sigma^{\cdot}, 0, u^{\cdot}_{ij}), & l \in [0, u^{\cdot}_{ij}], \\ (1-z)\, f_{GPD}(l;u^{\cdot}_{ij}, \xi^{\cdot}_{ij}, \beta^{\cdot}_{ij}), & l \in (u^{\cdot}_{ij}, \infty), \end{cases} \\
\mu^{\cdot}_{ij} &= \alpha_{\mu,\cdot} + \sum_k f_{\mu,\cdot,k}(x_{jk}) + g_{\mu,\cdot}(t_i), \\
\xi^{\cdot}_{ij} &= \alpha_{\xi,\cdot} + \sum_k f_{\xi,\cdot,k}(x_{jk}) + g_{\xi,\cdot}(t_i), \\
\beta^{\cdot}_{ij} &= f_{\beta,\cdot}(x_j, t_i),
\end{aligned} \tag{13}$$

where $TruncLN(\mu, \sigma, x_{\min}, x_{\max})$ denotes a truncated log-normal distribution on the interval $[x_{\min}, x_{\max}]$ and $GPD(u, \xi, \beta)$ denotes a generalized Pareto distribution with

[16] Note that when fitting a spliced severity distribution as below, in order to apply established fitting procedures, one would usually select a global, non-covariate-dependent threshold u and fit each distribution onto the data that fall into the "globally" specified regions. As we do not address the question of model fitting here, we stick to the more general formulation, as it is interesting to assume that depending on the covariates, the classification of a severity as *extreme* should start at different levels.

location u, shape ξ, and scale β.[17] As for the idiosyncratic frequency modelling illustrated in Sect. 4.2.1, the sum might run over different subsets of covariates for different incident types (see e.g. Eq. (3)).

Next, we combine the concepts for frequency and severity modelling to study some questions that arise from an actuarial viewpoint.

4.4 Insurance pricing and risk measurement

Recall that we take the perspective of an insurer, whose portfolio consists of K firms exposed to cyber losses whose frequency and severity are modelled as detailed in Sects. 4.2 and 4.3. Questions of interest for the insurer when setting up a portfolio of cyber insurance policies typically include:

1. Contract design (deductibles, cover limits, coverage period);
2. Pricing of individual policies given an applicant's characteristics;
3. Estimation and quantification of the portfolio risk.

In this work, we do not elaborate in detail on the first question and for now assume no deductible, no cover limit, and a standard policy duration of one year. These assumptions imply that for each incident, the loss suffered by the insured firm and the claim size faced by the insurer are equal and the terms will be used interchangeably. To study the latter two questions, the *total claim amount process* is denoted

$$L(t) = \sum_{i=1}^{N(t)} Y_i, \quad t \geq 0,$$

where it is assumed that the claim number process $(N(t))_{t\geq 0}$ is independent of the iid., a.s. positive, claim size sequence $\{Y_i\}_{i\in\mathbb{N}}$. We restrict our focus to the case of fixed $T > 0$, i.e. instead of studying the process $(L(t))_{t\geq 0}$, study the random variable $L(T)$. In general, it is very hard to make statements about the exact distribution of L and one has to resort to Monte Carlo methods or, if applicable, a numerical routine like the *Panjer recursion*.

In our context, the loss for a firm $j \in \{1, \dots, K\}$ up to time $T > 0$ from one type of cyber incidents (e.g. data breaches) can be expressed as

$$L_j^{\cdot}(T) = \sum_{i=1}^{N_j^{\cdot}(T)} L_i^{\cdot,(j)},$$

[17] Note that when fitting a GPD with covariate-dependent parameters using the method developed in Ref. [57], an orthogonal reparametrization $(\xi(x_j, t), \nu(x_j, t)) := (\xi(x_j, t), \log(\beta(x_j, t)(1 + \xi(x_j, t))))$ is chosen. The resulting MLE $\hat{\nu}$ can be transformed back directly to an estimator $\hat{\beta}$, but the dependence of β on the covariates does then not follow a GAM structure anymore. Therefore, a more general functional relationship is stated above.

where $N_j^{\cdot}(T) \sim Poi\big(\Lambda_j^{\cdot}(T)\big)$ as given in Eq. (11) and $\{L_i^{\cdot,(j)}\}_{i\in\mathbb{N}} := \{Y^{\cdot,(j)}(t_i)\}_{i\in\mathbb{N}}$ with pdf. $f_{L_{ij}^{\cdot}}$ as given in Eq. (13). Note that in general the sequence $\{L_i^{\cdot,(j)}\}_{i\in\mathbb{N}}$ is not iid. (due to time-dependence). However, if we drop time-dependence, which in practice could mean assuming constant severity distributions on one-year intervals considered separately, $L_j^{\cdot}(T)$ is again compound Poisson, and the total cyber loss incurred by firm $j \in \{1,\dots,K\}$ is given by

$$L_j(T) = \sum_{i=1}^{N_j(T)} L_i^{(j)}, \text{ where } N_j(T) \sim Poi\big(\Lambda_j^{DB}(T) + \Lambda_j^{FR}(T) + \Lambda_j^{BI}(T)\big),$$

$$\text{and } F_{L_i^{(j)}} = \sum_{y\in\{DB,FR,BI\}} \frac{\Lambda_j^{y}(T)}{\Lambda_j^{DB}(T) + \Lambda_j^{FR}(T) + \Lambda_j^{BI}(T)} F_{L_i^{y,(j)}}.$$

The portfolio loss is simply given by the sum of (dependent) firm losses, i.e.

$$L^{\cdot}(T) = \sum_{j=1}^{K} L_j^{\cdot}(T) \text{ and } L(T) = \sum_{j=1}^{K} L_j(T).$$

Regarding the second question of finding a premium $\Pi(T)$ for an individual insurance policy on [0, T], we recall the well-known premium calculation principles listed below [86]. As it is often impossible to find the exact distributional properties of the total claim amount process, the ones based on the first two moments are popular in practice.

- Expected value principle: $\Pi_j(T) = (1+\rho)\mathbb{E}\big[L_j(T)\big]$, with safety loading $\rho > 0$.
- Standard deviation principle: $\Pi_j(T) = \mathbb{E}\big[L_j(T)\big] + \rho\sqrt{\mathbb{V}ar\big(L_j(T)\big)}$, where $\rho > 0$.
- Exponential principle: $\Pi_j(T) = \frac{1}{\gamma}\log\big(\mathbb{E}[e^{\gamma L_j(T)}]\big)$, with risk aversion $\gamma > 0$.

Concerning the question of quantifying the risk of the overall portfolio loss, the two most common tail risk measures are the *Value-at-Risk (VaR)* at a given confidence level $1-\alpha$ and, if applicable, the corresponding *Average Value-at-Risk (AVaR)*. Theoretically, for a positive loss r.v. L with cdf F_L, they are given by

$$\begin{aligned} VaR_{1-\alpha}(L) &:= \inf\big\{l \in \mathbb{R} : \mathbb{P}(L \le l) \ge 1-\alpha\big\} = F_L^{-1}(1-\alpha), \\ AVaR_{1-\alpha}(L) &:= \mathbb{E}\big[L \big| L \ge VaR_{1-\alpha}(L)\big] \overset{(*)}{=} \frac{1}{1-\alpha}\int_0^{1-\alpha} VaR_{\gamma}(L)\mathrm{d}\gamma, \end{aligned}$$

where F_L^{-1} denotes the generalized inverse of F_L and $(*)$ requires F_L to be continuous. Note that in cases with very heavy-tailed loss severities (as e.g. observed in some of the previous works on cyber risk), $AVaR(L)$ cannot be computed as it relies on L to have finite expectation.

5 An example of an actuarial application via a simulation study

The aim of the following section is to illustrate the application of the proposed modelling approach to pricing and risk measurement in an actuarial context. To this end, a fictitious insurance portfolio is constructed and parameters for the frequency and severity distributions as given in the previous sections are proposed based on previous academic literature and expert judgement. Based on the resulting simulated portfolio loss distribution, the effect of interdependent losses and the introduction of cover limits is highlighted. Due to the scarcity of available empirical data, the parameters and model assumptions could not yet be fit to (resp. challenged on) a real dataset; this remains an important task for future research.

5.1 Portfolio composition and company covariates

We first construct a (fictitious) insurance portfolio consisting of $K = 50$ firms from $B = 6$ sectors, all details are listed in Table 3. A bigger portfolio, which is used in our simulation study, is then obtained by copying each firm 10 times with IT security levels varying from 0.05 to 0.95 (stepsize 0.1). This enables us to compare the results of the entire portfolio ($K = 500$) with sub-portfolios ($K = 50$) of different security level (denoted sub-portfolio 1 – 10), and for each individual firm with varying security level. Table 4 gives an overview of the relative and absolute frequencies for each covariate in each sub-portfolio.

5.2 Frequency distribution

We require our simulation to adhere to the following stylized facts (F1)–(F5) for the frequency of idiosyncratic incidents:

(F1) Consider a $T = 5$-year observation period, during which the frequency increases by around 67% [3]. The increase is realized in yearly (log-linear) steps; within each year the frequency is assumed constant.

(F2) During the first year ($t \in [0, 1)$) and for baseline covariate levels $s_j = d_j = nsup_j = 1, c_j = 0.5$, the incident (loss) probability is 0.01 (this is a conservative estimate).

(F3) Incidents are distributed into 25% DBs [56], 25% BIs, and 50% FR (or *other*).

(F4) An increase of either of the categorical covariates s_j, d_j, and $nsup_j$ by one (two) level(s) from the baseline implies an increase of the incident rate by 10% (20%).

(F5) Assume a log-linear influence of the security level such that increasing it to the maximum ($c_j = 1$) yields a halved rate of cyber incidents (and thus lowering it to the minimum ($c_j = 0$) leads to a doubled rate), compared to the baseline.

These assumptions imply the parameters for the covariate-dependent rates of idiosyncratic incidents (c.f. Eq. (2)) in the upper panel of Table 5.[18] For systemic events, we follow Sect. 4.2.2, where the assumed parameters for the ground process (c.f. Eq. (4)) and the mark distribution are given in the lower panel of Table 5. Although it is difficult to make assumptions, as none of the existing studies explicitly distinguish systemic events, Table 5 reflects the following simplifying assumptions (F6), (F7) (as before, we do not have any information to justify any more complex assumptions):

(F6) The mark distribution is equal for DB, BI, and FR. Sector-specific events are (discretely) uniformly distributed over all sectors.

(F7) The number of incidents from systemic events is similar to the number of idiosyncratic incidents for baseline covariate levels, which implies a doubled overall incident frequency (and a 50% increased loss frequency).

5.3 Severity distribution

For this study, we deviate from the very high mean (resp. median) severity estimates given in the existing literature (several million US$ for a single incident) for two reasons: First, it is reasonable that events listed in public databases exhibit much higher losses than the average *daily-life* cyber incident that goes unnoticed by the public and second, insurance policies currently offered on the market (especially policies for SMEs) usually have cover limits of up to 5 million US$, therefore it would not be reasonable to assume mean claim severities that already exhaust the policy limit.[19] Recall that this study is intended as a prototype to show the general behaviour of the model; absolute numbers given should not be interpreted as representative of

[18] To illustrate how these parameters relate to the assumptions, take the example of (F4): The increase of the idiosyncratic rate of some type of incident when increasing a categorical covariate by one level from the benchmark (where the benchmark is represented by the intercept) is given by $\lambda^{\cdot,idio}((x_{j1}, 2, x_{j3}, x_{j4}, x_{j5}), t)/\lambda^{\cdot,idio}((x_{j1}, 1, x_{j3}, x_{j4}, x_{j5}), t) = \exp(f_{\lambda,\cdot,2}(2))$. Equating this ratio to 1.1, i.e. assuming a c.p. 10% increase, yields $f_{\lambda,\cdot,2}(2) = 0.095$. Likewise, equating $\lambda^{\cdot,idio}((x_{j1}, 3, x_{j3}, x_{j4}, x_{j5}), t)/\lambda^{\cdot,idio}((x_{j1}, 1, x_{j3}, x_{j4}, x_{j5}), t) = \exp(f_{\lambda,\cdot,2}(3))$ to 1.2, i.e. assuming a c.p. 20% increase, yields $f_{\lambda,\cdot,2}(3) = 0.18$.

[19] Note that as the existing studies do not state whether the recorded cyber losses were fully or partly insured, it is not possible to make statements about the relationship between those losses and the size of potentially corresponding insurance claims.

a real-world portfolio. The following assumptions (S1)–(S7) lead to the choice of parameters given in Table 6:

(S1) During the first year and for baseline covariate levels, for all types of incidents the expected claim size of the underlying log-normal distribution is given by $\mathbb{E}[\widetilde{L}_{ij}] = 50$.

(S2) The standard deviation of the underlying log-normal cost distribution is constant and consistent with the results for (negligent) data breaches in Refs. [49] and [72].

(S3) The expected claim size $\mathbb{E}[\widetilde{L}_{ij}]$ increases by 10% (20%) for a one (two) level increase of either s_j or d_j relative to the benchmark. The influence of c_j on $\mathbb{E}[\widetilde{L}_{ij}]$ is log-linear, where $c_j = 1$ results in a halved expected claim size.

(S4) Over the $T = 5$-year observation period, $\mathbb{E}[\widetilde{L}_{ij}]$ increases (in yearly log-linear steps) by 60%.

(S5) For *large claims*, the shape parameter ξ of the GPD is constant and close to 1 [56] to model heavy-tailed behaviour while avoiding a switch from a finite-mean to an infinite-mean scenario.

(S6) The expected threshold exceedance (relative to the corresponding threshold, dependent on the underlying log-normal distribution) $\mathbb{E}[L_{ij} - u_{ij} \mid L_{ij} > u_{ij}]/u_{ij} = \beta_{ij}\,(u_{ij}(1 - \xi_{ij}))^{-1}$ equals 0.5 for baseline covariate levels, i.e. the expected size of a claim exceeding the threshold is given by 1.5 times the threshold.

(S7) The same assumptions regarding covariate- and time-dependence as for small claims are made, referring to the expected relative threshold exceedance (e.g. a one-level increase of s_j leads to a 10% increase) instead of the expected claim size. In this case, the influence of c_j is linear and such that $c_j = 1$ results in a halved expected relative threshold exceedance.[20]

5.4 Results of the simulation study

The following results are based on 50.000 simulation runs on a grid of 5 years, reported values refer to the first year unless stated otherwise. For each run, the arrival times of idiosyncratic incidents (at each firm) and systemic events are generated using the rates in Eqs. (2) and (4), respectively. For each systemic event, the affected subset S_i is generated as described in ($A4$) using r.v. G_i, B_i, and Z_{ij} from their respective distributions. Furthermore, m_i is drawn and the set S_i^* deduced from the realisations of S_i and m_i. This

[20] Assumptions (S6) and (S7) result in equations for β of the type $\beta^{\cdot}_{ij} = u^{\cdot}_{ij}\,(1 - \xi^{\cdot})\,(\alpha_{\beta,\cdot} + \sum f_{\beta,\cdot,k}(x_{jk}) + g_{\beta,\cdot}(t))$ with coefficients given in Table 6 which do not strictly fit into the framework of [57] for fitting a covariate-dependent GPD. When calibrating the model to data, it is not required to make any such assumption. Note, however, that due to the reparametrization in the framework of [57], the covariate dependence of β is not intuitive. Therefore, we stick to intuitively interpretable assumptions.

Table 3 List and modelling framework of company covariates to be included in a cyber risk model

Covariate	Abbreviation	Type	Scope	Information availability	Comment
Industry sector	*b*	Categorical	FI: finance and insurance BR: businesses (retail) HC: healthcare EDU: education GOV: government and military MAN: manufacturing	Public data	All but the last sector are covered in the PRC dataset and therefore used in the works using that data (e.g. Refs. [50–52]) Sector found as relevant covariate in Refs. [3, 5, 51, 52, 54, 55].
Size	*s*	Ordinal	1 Small 2 Medium 3 Large	Public data or questionnaire	Usually, for the size determination of an enterprise, revenue and number of employees should be used jointly. For details, see Table 11 in Online Appendix A.7. Size found as relevant covariate in Refs. [53–55].
Data	*d*	Ordinal	1 Low risk 2 Medium risk 3 High risk	Self-report via questionnaire, otherwise approximate using public data about industry sector and size	Data found as relevant covariate in Refs. [52, 55, 72]. Usually considered for insurance pricing [15]. Risk classification should take into account number of records and type of data (e.g. PII, PHI, credit card data is more sensitive), see Table 12 in Online Appendix A.7.

Table 3 (continued)

Covariate	Abbreviation	Type	Scope	Information availability	Comment
IT security level	c	Numerical	$[c_{min}, c_{max}] \overset{\text{w.l.o.g.}}{=} [0, 1]$	Self-report via questionnaire or e.g. scrutiny by a service provider hired by the insurer	Relevance is clear, but e.g. emphasized in Refs. [5, 60]. Usually taken into account for insurance pricing [15].
Number of suppliers	*nsup*	Ordinal	1 Low 2 Medium 3 High	Hard to elicit, could be estimated from industry sector and size, or via questionnaire	Depends on sector and number of employees. For details, see Table 13 in Online Appendix A.7. Relevant for exposure to supplier attacks.

yields an overall number of losses and incidents for each firm on each step of the grid, such that the corresponding severities can be drawn from the appropriate (time- and covariate-dependent) distribution.

5.4.1 Cumulative loss distribution

First, we examine the number of incidents/losses and the distribution of cumulative losses in the full portfolio ($K = 500$) in Fig. 2a and b. We compare the case where only actual losses are counted with the case where all incidents are counted (i.e. the *worst case* where all incidents cause a loss). At first glance, the two cases appear to be surprisingly similar. Notice, however, that due to the assumptions above, for most firms the rate of idiosyncratic incidents outweighs the rate of incidents from systemic events. Furthermore, the lower the security level for a given firm, the higher its contribution to the overall number of incidents, and simultaneously the lower the effect of distinguishing losses and incidents. Conversely, the higher the security level of a given firm, the less likely it is to be affected at all. Therefore, when only few cases are registered at all, these cases are likely to have occurred at firms with low security and are therefore unlikely to be filtered. The cases where most filtering occurs are large systemic events whose effect is clearly reduced (consider the range around [45, 65] on the x-axis of Fig. 2b). Of course, this translates analogously to Fig. 2a, where particularly the tail of the distribution is altered (x-axis-range around [2500, 4000] in Fig. 2a). In this case, it additionally has to be kept in mind that incidents at well-protected firms—which are mostly filtered—are assumed to typically cause below-average losses. In both figures, one observes the difference in mean between counting losses and incidents, and that the mean is shifted clearly to the right from the mode of the body of the distribution. As expected from the assumptions above, we observe a shift of the cumulative loss distribution to the right as time progresses. To corroborate the simulation results, we generate 50.000 samples of incident/loss numbers following Proposition 1 and Eq. (12) and compare them in Fig. 2d. The simulation via Eq. (12) is much faster, but cannot be directly used to generate the cumulative loss distribution, as only samples of the *total number* of incidents/losses are drawn without information as to which firms they affect (and severity differs between firms).

Furthermore, we compare the cumulative loss distribution for selected sub-portfolios in Fig. 2c, taking into account only simulation runs where a non-zero loss has been observed. As to be expected, we observe a shift of the body and tail of the loss distribution to the left as the security level increases. Understanding the cumulative loss distribution—especially in the tail—is particularly interesting in the context of reinsurance, where common contract design involves so-called *excess-of-loss* reinsurance, meaning that (portfolio) losses exceeding a pre-specified limit are ceded. For this case, an accurate understanding of the portfolio loss distribution and its tail

Table 4 Occurrence frequencies of covariate values in the toy portfolio

Covariate	Scope	Frequency
Sector b_j	FI: finance and insurance	0.30 (15)
	HC: healthcare	0.30 (15)
	BR: businesses (retail)	0.10 (5)
	EDU: education	0.10 (5)
	GOV: government and military	0.10 (5)
	MAN: manufacturing	0.10 (5)
Size s_j	1 Small	0.60 (30)
	2 Medium	0.30 (15)
	3 Large	0.10 (5)
Data d_j	1 Low risk	0.20 (10)
	2 Medium risk	0.28 (14)
	3 High risk	0.52 (26)
Number of suppliers $nsup_j$	1 Low	0.74 (37)
	2 Medium	0.20 (10)
	3 High	0.06 (3)

This dataset is copied ten times with varying IT security level ranging from 0.05 to 0.95

Table 5 Chosen parameter assumptions for frequencies (based on (F1)–(F7))

Idiosyncratic incidents		
Intercept	$(\alpha_{DB}, \alpha_{FR}, \alpha_{BI})$	$(-6, -5.3, -6)$
Data factor levels	$f_{DB,3}(x_{j3})$	$(0, 0.095, 0.18)$
Size factor levels	$f_{FR,2}(x_{j2}), f_{BI,2}(x_{j2})$	$(0, 0.095, 0.18)$
Supplier factor levels	$f_{DB,5}(x_{j5}), f_{FR,5}(x_{j5}), f_{BI,5}(x_{j5})$	$(0, 0.095, 0.18)$
IT security dependence	$f_{DB,4}(x_{j4}), f_{BI,4}(x_{j4})$	$1.39\ (0.5 - x_{j4})$
Time dependence	$g_{\lambda^{DB,idio}}(t), g_{\lambda^{FR,idio}}(t), g_{\lambda^{BI,idio}}(t)$	$0.128\ \lfloor t \rfloor$
Ground process of systemic events		
$\lambda^{DB,g}(t) = \exp(g_{\lambda^{DB,g}}(t))$	$\exp(-3.28 + 0.128\ \lfloor t \rfloor)$	
$\lambda^{FR,g}(t) = \exp(g_{\lambda^{FR,g}}(t))$	$\exp(-2.59 + 0.128\ \lfloor t \rfloor)$	
$\lambda^{BI,g}(t) = \exp(g_{\lambda^{BI,g}}(t))$	$\exp(-3.28 + 0.128\ \lfloor t \rfloor)$	
Distribution of S_i		
(p_G, p_{gen}, p_{sec})	$(0.5, 0.1, 0.2)$	
Sector distribution	$B_i \sim Unif\{1, \dots, 6\}$, i.e. $p_b = \frac{1}{6}\ \forall b \in \{1, \dots, B\}$	

is clearly essential. Apart from these considerations, insurers are mostly concerned with the pricing of individual policies. This is addressed next.

Table 6 Chosen parameter assumptions for severities (based on (S1)–(S7))

μ		
Intercept	$\alpha_{\mu,\cdot}$	3.91
Data factor levels	$f_{\mu,DB,3}(x_{j3})$	(0, 0.095, 0.18)
Size factor levels	$f_{\mu,FR,2}(x_{j2}),f_{\mu,BI,2}(x_{j2})$	(0, 0.095, 0.18)
IT security dependence	$f_{\mu,\cdot,4}(x_{j4})$	$1.39\ (0.5 - x_{j4})$
Time dependence	$g_{\mu,\cdot}(t)$	$0.1175\ \lfloor t \rfloor$
σ	$\sigma^{\cdot}$	0.076
ξ	$\alpha_{\xi,\cdot}$	0.9
β		
Intercept	$\alpha_{\beta,\cdot}$	0.5
Data factor levels	$f_{\beta,DB,3}(x_{j3})$	(0, 0.05, 0.1)
Size factor levels	$f_{\beta,FR,2}(x_{j2}),f_{\beta,BI,2}(x_{j2})$	(0, 0.05, 0.1)
IT security dependence	$f_{\beta,\cdot,4}(x_{j4})$	$0.5\ (0.5 - x_{j4})$
Time dependence	$g_{\beta,\cdot}(t)$	$(0, 0.063, 0.133, 0.211, 0.3)\ \mathbb{1}_{\{\lfloor t \rfloor = i\}},\ i \in \{0, \ldots, 4\}$

5.4.2 Premium calculation

Based on the distribution of individual losses, we calculate the first-year premium based on the expected value principle given in Sect. 4.4.[21] Table 7 compares the following exemplary firms:

Firm 1: A small manufacturing business with low data and supplier risk and low IT security standards ($c = 0.15$).
Firm 2: A medium-sized company in the financial sector with medium data and supplier risk and high IT security standards ($c = 0.85$).
Firm 3: A large health care provider with high data risk, medium supplier risk, and average IT security standards ($c = 0.55$).

The results show that the IT security level dominates the covariate-effect on the premium, which is in line with the assumptions. As in practice, the calculation of expected losses is typically based on historically recorded (rare!) losses, only very few firms with the exact same covariate combinations might be in the portfolio and therefore, the premium is rather calculated based on all losses within a class of firms considered *homogeneous*. New firms falling into the same class are then assigned the same premium. The quite difficult task is to find an appropriate way of partitioning firms into homogeneous groups. If we partition firms according to their IT security level and calculate their premium by taking into account all firms with the same level, we obtain the results shown in Fig. 3.[22] As to be expected, the premium

[21] Note that for the chosen severity parameters, only the first moment exists ($0.5 < \xi < 1$). This prohibits the use of the exponential and standard deviation principle. We will remedy this by introducing cover limits later.

[22] *Theoretical* premiums in this figure refer to the premium that would be assigned to each firm if the expected sub-portfolio loss (the sum of the expected single losses) was allocated evenly among all firms in the sub-portfolio. This is analogous to the *simulated* approach of pricing each firm equally based on

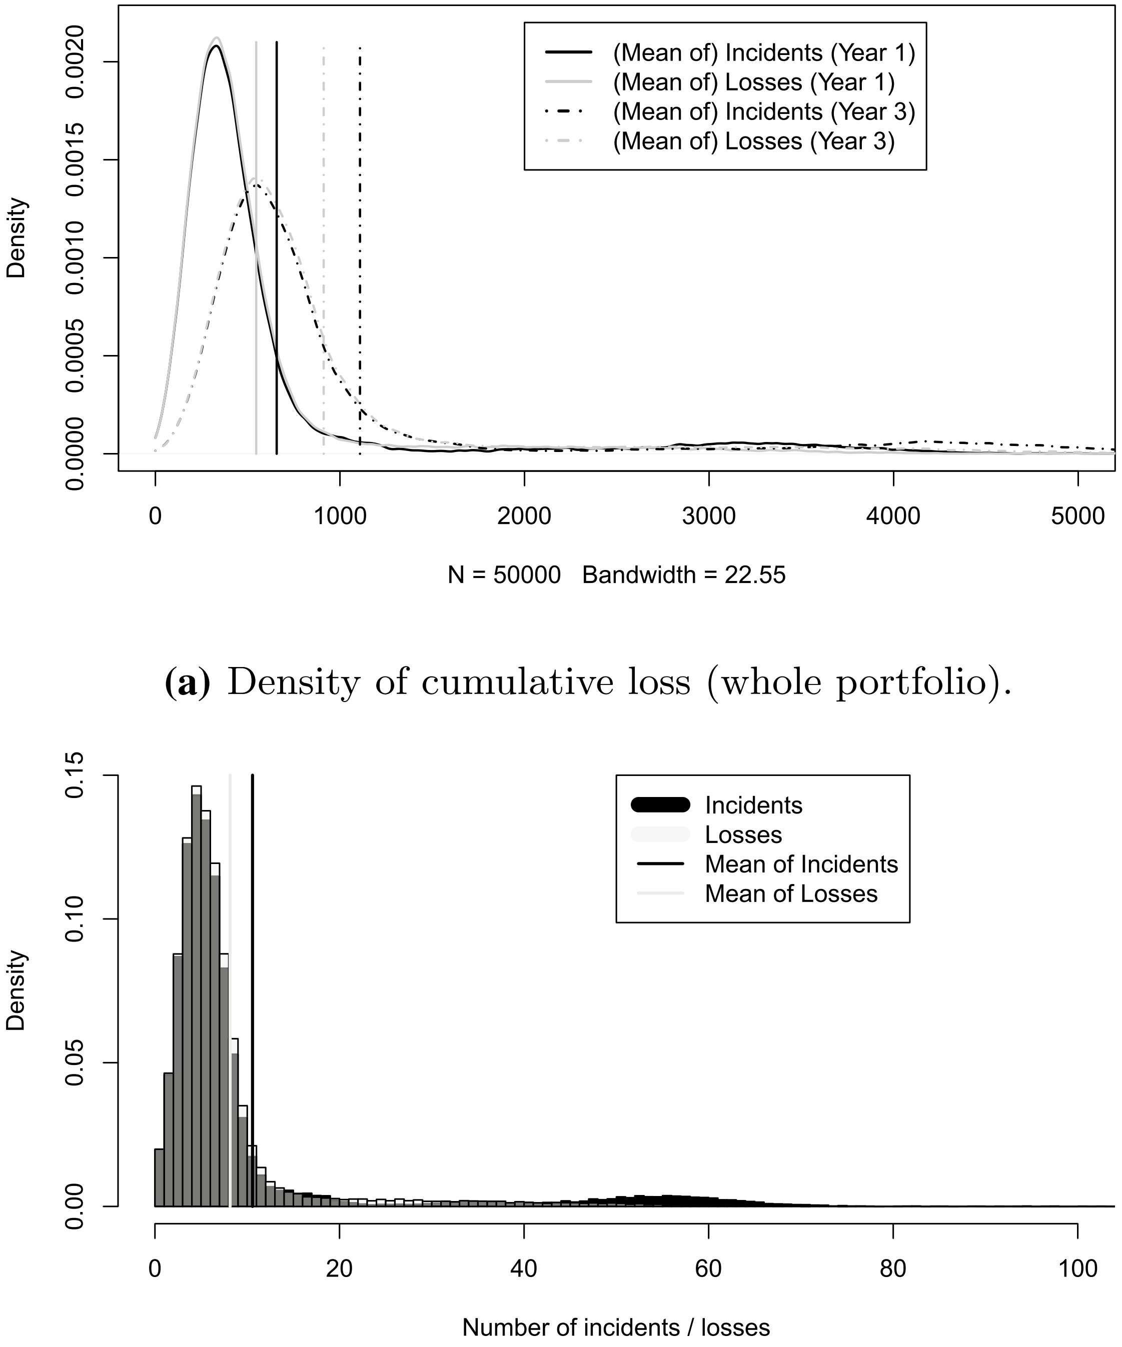

(a) Density of cumulative loss (whole portfolio).

(b) Histogram of incident / loss numbers.

Fig. 2 In panels (**a**) and (**b**), for the entire portfolio ($K = 500$), the number of incidents and losses and the distribution of the cumulative portfolio loss for 50.000 runs is compared. In panel (**c**), the cumulative loss distributions for three sub-portfolios with different security levels, namely 0.05 (Portfolio 1), 0.45 (Portfolio 5), and 0.95 (Portfolio 10) are shown; here, only runs with non-zero recorded loss are taken into account, causing the sample size to vary between portfolios as to be expected. In panel (**d**), incident numbers as in Eq. (12) are simulated such that one can observe the similarity to panel (**b**)

Footnote 22 (continued)
the loss history of the—assumed homogeneous—portfolio. Combining the two "extremes" of considering only individual loss experience and only loss experience from a homogeneous group lies at the heart of *credibility theory* approaches and will not be addressed here.

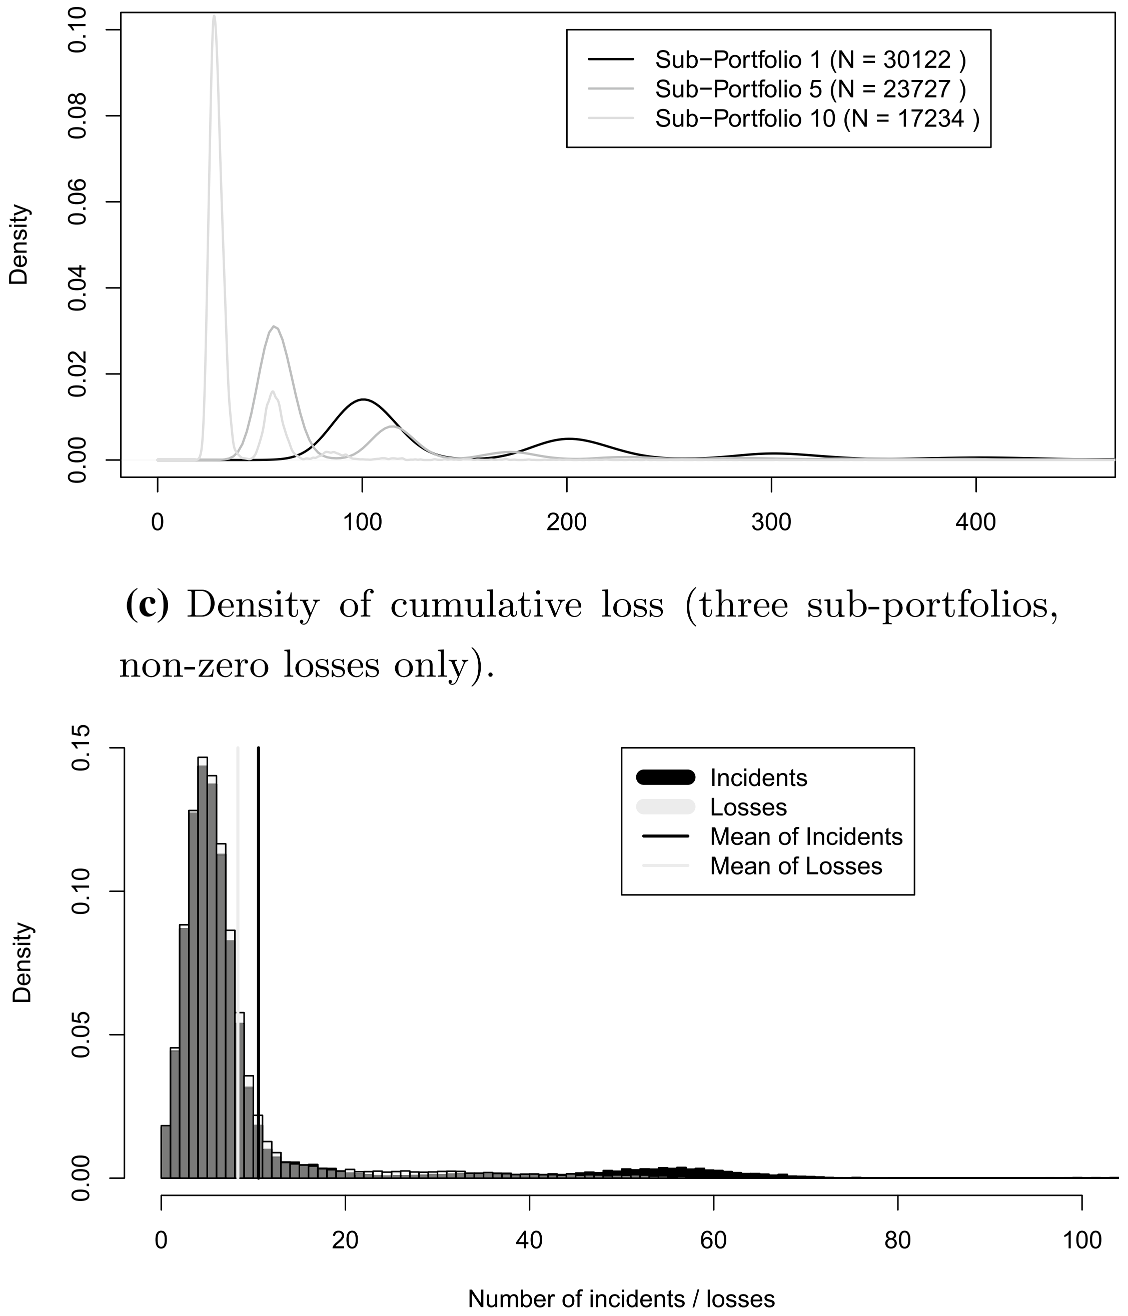

(c) Density of cumulative loss (three sub-portfolios, non-zero losses only).

(d) Histogram of incident / loss numbers (based on (5), (6), and (12)).

Fig. 2 (continued)

decreases with increasing security level, while the difference between incidents and losses increases. As an alternative to (and validation for) the Monte Carlo simulation, we have furthermore implemented a *Panjer recursion scheme* using a discretized version of the severity distribution; the results are given in Online Appendix A.5 and corroborate the ones given here.

Table 7 Comparison of first-year cyber insurance premium for three selected firms, simulated numbers based on 50.000 runs

Premium based on expected value principle ($\rho = 0.2$)

	Based on losses		Based on incidents	
	Theoretical	Simulated	Theoretical	Simulated
Firm 1	2.1665	2.0814	2.3174	2.2338
Firm 2	0.4610	0.4451	0.8107	0.7746
Firm 3	1.1777	1.1732	1.5557	1.5164

The difference between losses and incidents represents a reduction that can purely be achieved through enhanced security, as more incidents from systemic events are filtered. Note that enhanced security also decreases the frequency of idiosyncratic incidents; this is reflected in both cases

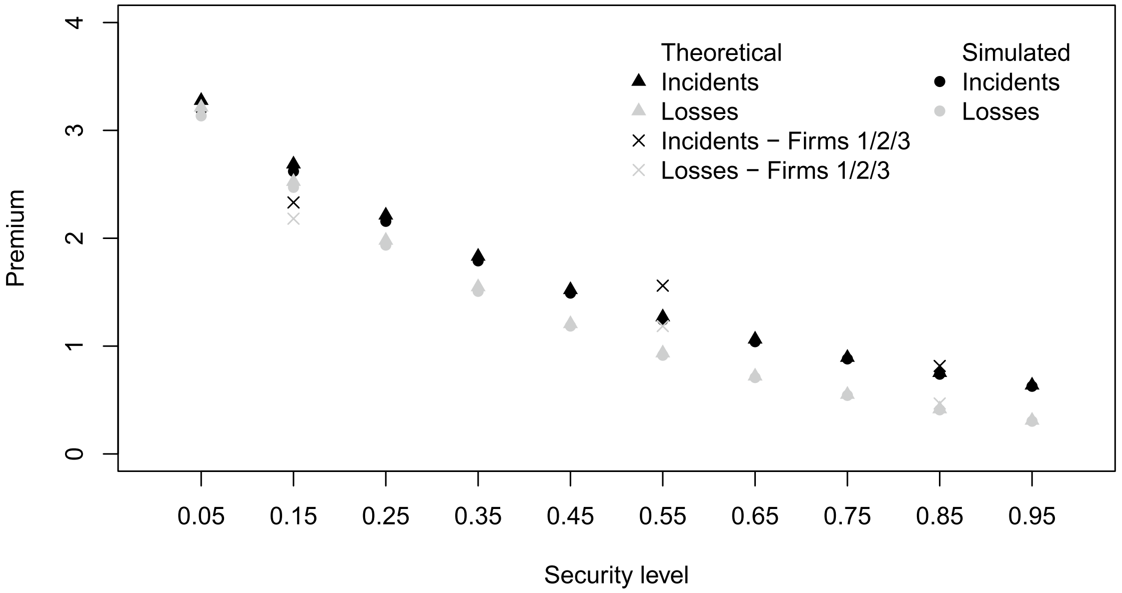

Fig. 3 We compare the premium (with loading 0.2 as above) that would be assigned to firms if they were grouped according to their IT security level. We observe that simulated values are now very close to theoretical ones, as they depend on the loss history of a sub-portfolio of 50 firms, such that Monte Carlo noise is reduced (compared to Table 7). We furthermore compare the values for the single firms from Table 7 with the portfolio they would be grouped into, and observe that e.g. firm 1, when evaluated on its own, is slightly less risky than the average firm in sub-portfolio 2

5.4.3 Risk measurement on individual and portfolio level

We compare *VaR* and *AVaR* for the three firms described above and two sub-portfolios in Table 8, as well as for all sub-portfolios in Fig. 4a and b. The *historical* estimate refers to the sample quantile from the simulation data, i.e. for a realisation of losses $\mathbf{L} = (L_1, \dots, L_n)$, let $L_{(1)} < L_{(2)} < \dots < L_{(n)}$ denote the order statistics, then for a chosen level $(1-\alpha) \in \left(\frac{i-1}{n}, \frac{i}{n}\right]$, $VaR_{1-\alpha}$ and $AVaR_{1-\alpha}$ are estimated as their empirical counterparts

$$\widehat{VaR}_{1-\alpha}(\mathbf{L}) = \hat{F}_L^{-1}(1-\alpha) = L_{(i)}, \quad \widehat{AVaR}_{1-\alpha}(\mathbf{L}) = \frac{1}{n-i+1}\sum_{j=i}^{n} L_{(j)}.$$

The *POT* estimate assumes that for a large threshold u, the excesses are distributed according to a generalized Pareto distribution $GPD(u, \xi, \beta)$, and thus $VaR_{1-\alpha}$ and $AVaR_{1-\alpha}$ can be estimated as (see, e.g. Ref. [57])

$$\widehat{VaR}_{1-\alpha}(\mathbf{L}) = u + \frac{\hat{\beta}}{\hat{\xi}}\left(\left(\frac{\alpha}{\frac{n'}{n}}\right)^{-\hat{\xi}} - 1\right), \quad \widehat{AVaR}_{1-\alpha}(\mathbf{L}) = \begin{cases} \frac{\widehat{VaR}_{1-\alpha}(\mathbf{L})+\hat{\beta}-\hat{\xi}u}{1-\hat{\xi}}, & \text{if} \quad \hat{\xi} \in (0,1), \\ \infty, & \text{if} \quad \hat{\xi} \geq 1, \end{cases}$$

where $\hat{\beta}$ and $\hat{\xi}$ are the parameter estimates of the scale and shape of the GPD given the data $\mathbf{L}$ and n' is the number of threshold exceedances. As to be expected, both *VaR* and *AVaR* decrease with increasing security level, while the reduction when considering only losses instead of all incidents is more substantial. Note again that

Table 8 Comparison of $VaR_{0.995}$ and $AVaR_{0.995}$ for the three selected firms and two selected sub-portfolios

Risk measures	$VaR_{0.995}$				$AVaR_{0.995}$			
	Losses		Incidents		Losses		Incidents	
	Hist	POT	Hist	POT	Hist	POT	Hist	POT
Firm 1	86.46	85.86	86.73	86.27	97.18	95.73	99.40	97.57
Firm 2	34.01	33.31	35.54	35.23	36.97	37.74	39.13	39.06
Firm 3	58.28	57.68	59.01	58.81	64.61	64.23	66.96	66.22
Portfolio 1	1056.01	1041.18	1067.52	1054.67	1408.11	1379.91	1423.35	1397.02
Portfolio 6	409.87	407.52	496.96	493.88	532.67	528.27	637.25	629.33

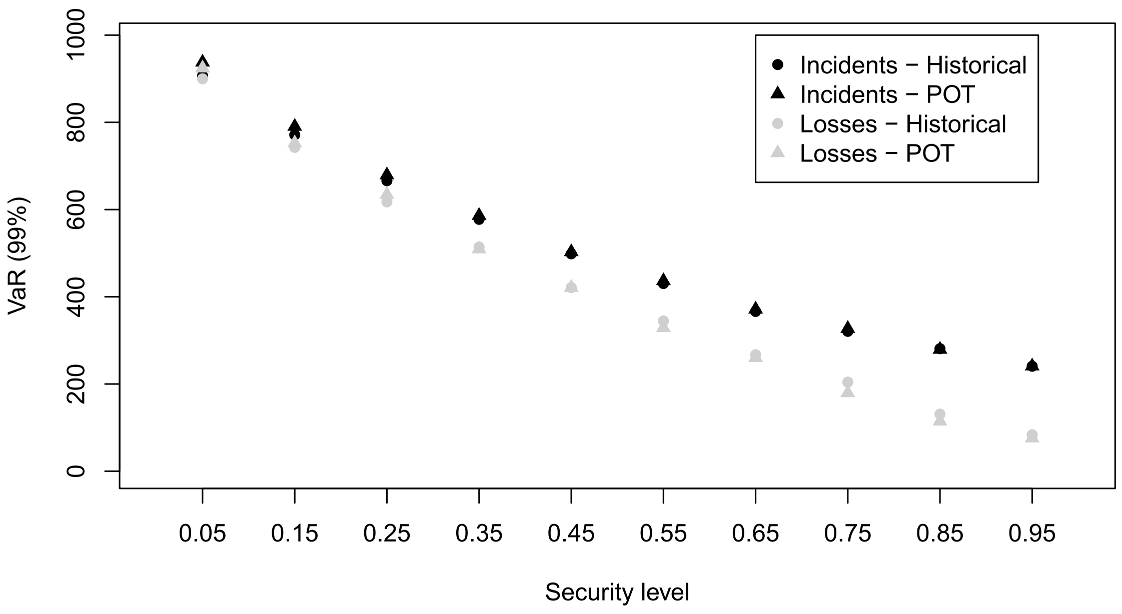

(a) $VaR_{0.99}$

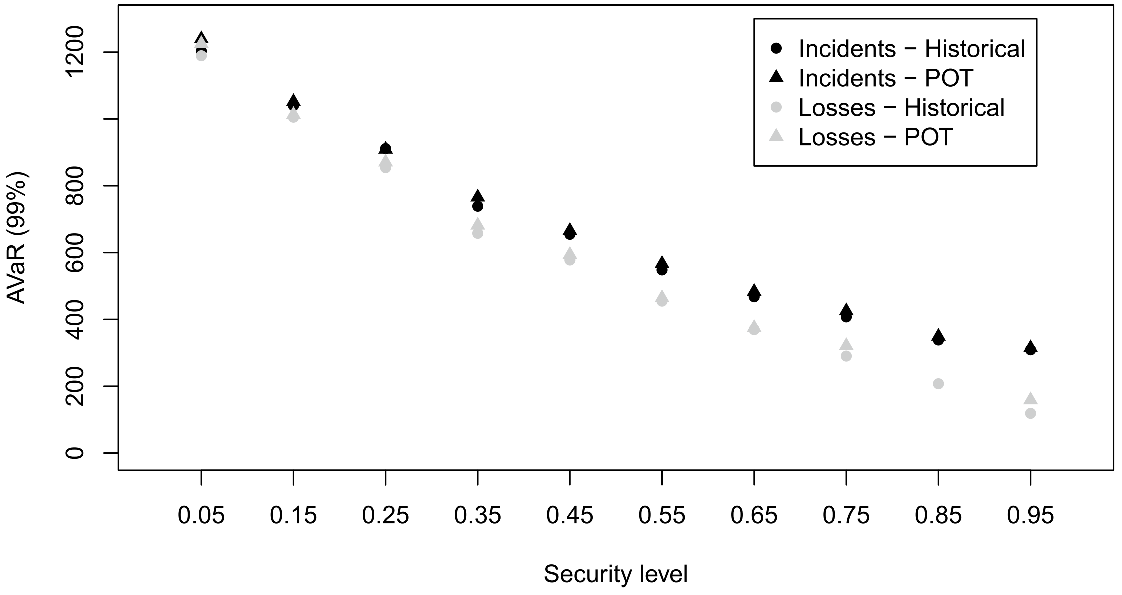

(b) $AVaR_{0.99}$

Fig. 4 Comparison of $VaR_{0.99}$ and $AVaR_{0.99}$ for all sub-portfolios

for individual firms, the numbers are based on their own loss history only and should be interpreted with care.

5.5 How relevant is accumulation risk?

We have repeatedly stressed the distinction between idiosyncratic incidents and systemic events and emphasized that the latter can lead to accumulation risk (re-)

insurers should be particularly worried about. One might now question whether the effect of including systemic events on the loss distribution warrants such a more complicated model. In order to answer this question (spoiler alert: yes!), we compare the results above with the results of a model that assumes the same marginal frequency as before for each firm, but assumes all incidents to be idiosyncratic, i.e. to occur independently from other firms. Intuitively, this should lead to the same premium for each individual contract, but decrease portfolio risk. Following Eq. (11), the overall number of incidents $\bar{N}_j^{\cdot}(T)$ at each firm $j \in \{1, \dots, K\}$ is generated using two independent Poisson r.v.

$$\bar{N}_j^{\cdot}(T) = \underbrace{N_j^{\cdot,idio}(T)}_{\sim Poi\left(\Lambda_j^{\cdot,idio}(T)\right)} + \underbrace{\bar{N}_j^{\cdot,syst}(T)}_{\sim Poi\left(\tilde{p}(b_j)\Lambda_j^{\cdot,g}(T)\right)}$$

independently from all other firms. To be able to compare the two cases in each run, the number of losses $N_j^{\cdot}(T)$ at each firm $j \in \{1, \dots, K\}$ is then generated as

$$N_j^{\cdot}(T) = N_j^{\cdot,idio}(T) + \underbrace{\bar{N}_j^{\cdot,syst}(T)}_{\sim Binom\left(N_j^{\cdot,syst}(T), \bar{F}_M(c_j)\right)} \quad .$$

With the severity distributions remaining unchanged, an analogous simulation study as above is conducted. Again, we first examine the overall distribution of the cumulative portfolio loss and number of incidents and losses in Fig. 5a and b, respectively. The difference to Fig. 2a and b is immediately evident:

- The visible heavy tails for both incident numbers and cumulative losses have vanished; thus it can be assumed they have been caused by systemic events with many firms affected simultaneously.
- In particular, the highest observed number of losses has decreased to around 17% of its previous value in both considered years, while mean losses and mean numbers of incidents/losses have stayed unaffected.
- The difference between incidents and losses is more directly visible, as in the independence case the body of the cumulative loss distribution is directly affected. This is because individual incidents are now filtered instead of the filtering impacting only systemic events, whose occurrence mostly alters the tail of the distribution.

From these findings, we conclude that incorporating systemic events into the model to capture potential accumulation risk is essential. We furthermore report $VaR_{0.99}$ and $AVaR_{0.99}$ for all sub-portfolios in Fig. 6a and b, respectively. Comparing them with Fig. 4a and b yields the same to-be-expected decreasing pattern as the security level increases, but the absolute values of the risk measures can be observed to have about halved. Perhaps it should rather be put vice versa: By including systemic events compared to complete independence, for the same expected overall number of incidents, the risk measures $VaR_{0.99}$ and $AVaR_{0.99}$ on sub-portfolio level *double*.

As the marginal frequency and severity for each firm remain unchanged, calculated premiums should not differ from the previous simulation study; this is corroborated in Online Appendix A.5.

We have mentioned before that very heavy-tailed loss severities characterize cyber risk, and previous studies typically suggest such heavy tails that any moments higher than first order do not exist (and in some cases tail parameter estimates even yield infinite-mean scenarios). Even a finite-mean, infinite-variance scenario (as above, with $0.5 < \xi < 1$) is cumbersome to deal with, as e.g. only premium calculations based on the first moment can be applied. "Luckily," in the insurance context, one typically does not deal with loss severities (without upper limit) directly, but rather with claim sizes, which are typically bounded from above by the introduction of a cover limit, a maximum amount the insurer is obliged to cover for each loss. The effects of this contract design feature are examined next.

5.6 Cyber policy design: the effect of cover limits

In practice, typical (primary) insurance contracts include a cover limit, as the insurer seeks to bound losses from single, extreme incidents. This, however, can lead to a supply-demand-mismatch: Insurers, still cautious of this new risk type, prefer relatively low cover limits (with a few exceptions, see the overview in Ref. [25]) that are sufficient to cover day-to-day cyber incidents, while many firms particularly seek protection for extreme scenarios such as a large data breach or long BI. [14, 58, 60] reported the non-existence of adequate cover limits as one reason for firms to refrain from purchasing cyber insurance.

Mathematically speaking, the introduction of a cover limit $\bar{M}$ corresponds to the truncation of the loss distribution, i.e. each $Y_i \in [0, \infty)$ is mapped to a claim size $\hat{Y}_i$ via $Y_i \mapsto \hat{Y}_i := \min\{Y_i, \bar{M}\} \in [0, \bar{M}]$. Note that we assume a limit on each loss; alternatives might be a limit on the total loss over the policy duration or a limit on the number of covered claims. Assuming, however, a realistically small claim frequency, this does not make a large difference, as cases of multiple losses happening at the same firm during a single policy year are extremely unlikely. Table 9 reports the probabilities of exceeding different cover limits for a large severity event and three different covariate combinations: the *baseline case* (year 1, $s = d = nsup = 1$, $c = 0.5$), the *lowest-risk case* in the portfolio (year 1, $s = d = nsup = 1$, $c = 0.95$), and the *highest-risk case* in the portfolio (year 5, $s = d = nsup = 3$, $c = 0.05$). To find the probability of an incoming claim to exceed the cover limit, we condition on observing a large claim event, i.e. in the notation of Sect. 4.3:

$$\mathbb{P}(L_{ij} > \bar{M}) = \underbrace{\mathbb{P}(L_{ij} > \bar{M} \mid L_{ij} > u_{ij})}_{\text{see Table 9}} \underbrace{\mathbb{P}(L_{ij} > u_{ij})}_{= 1-z \overset{e.g.}{=} 0.05}, \qquad \bar{M} > u_{ij}.$$

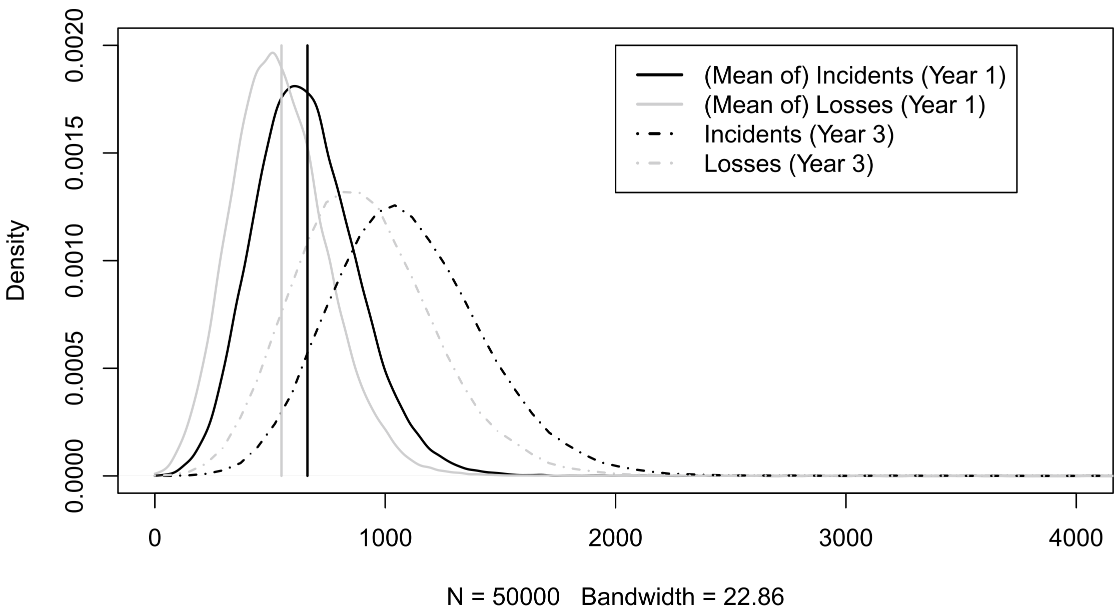

(a) Density of cumulative losses; Independence Case.

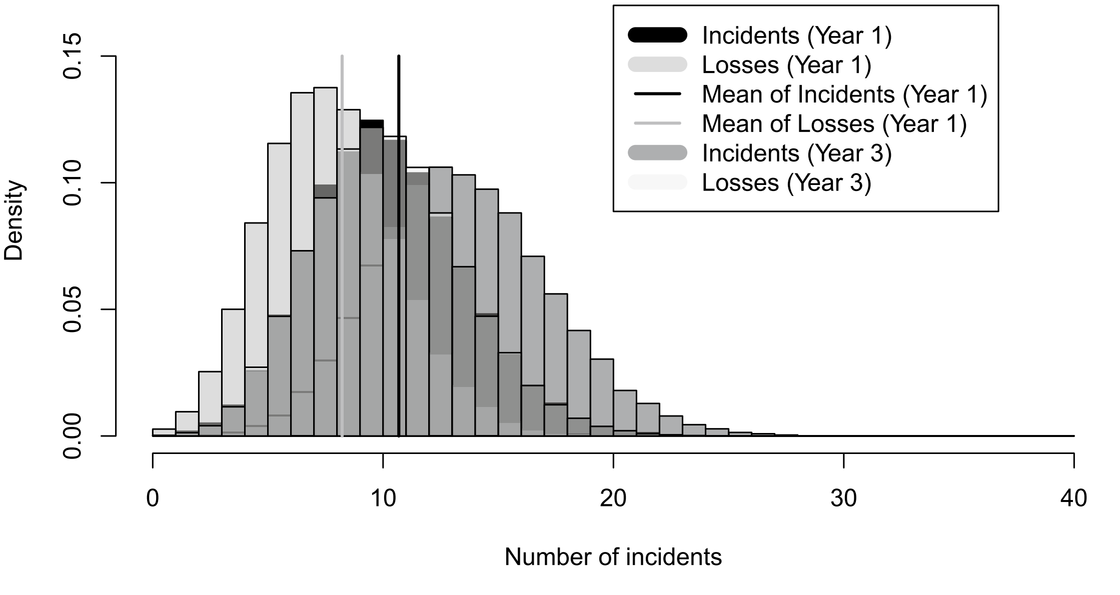

(b) Histogram of incident / loss numbers; Independence Case.

Fig. 5 For the entire portfolio ($K = 500$), the number of incidents/losses and the distribution of the cumulative portfolio loss for two different years is compared if incidents are assumed to arrive completely independently between firms

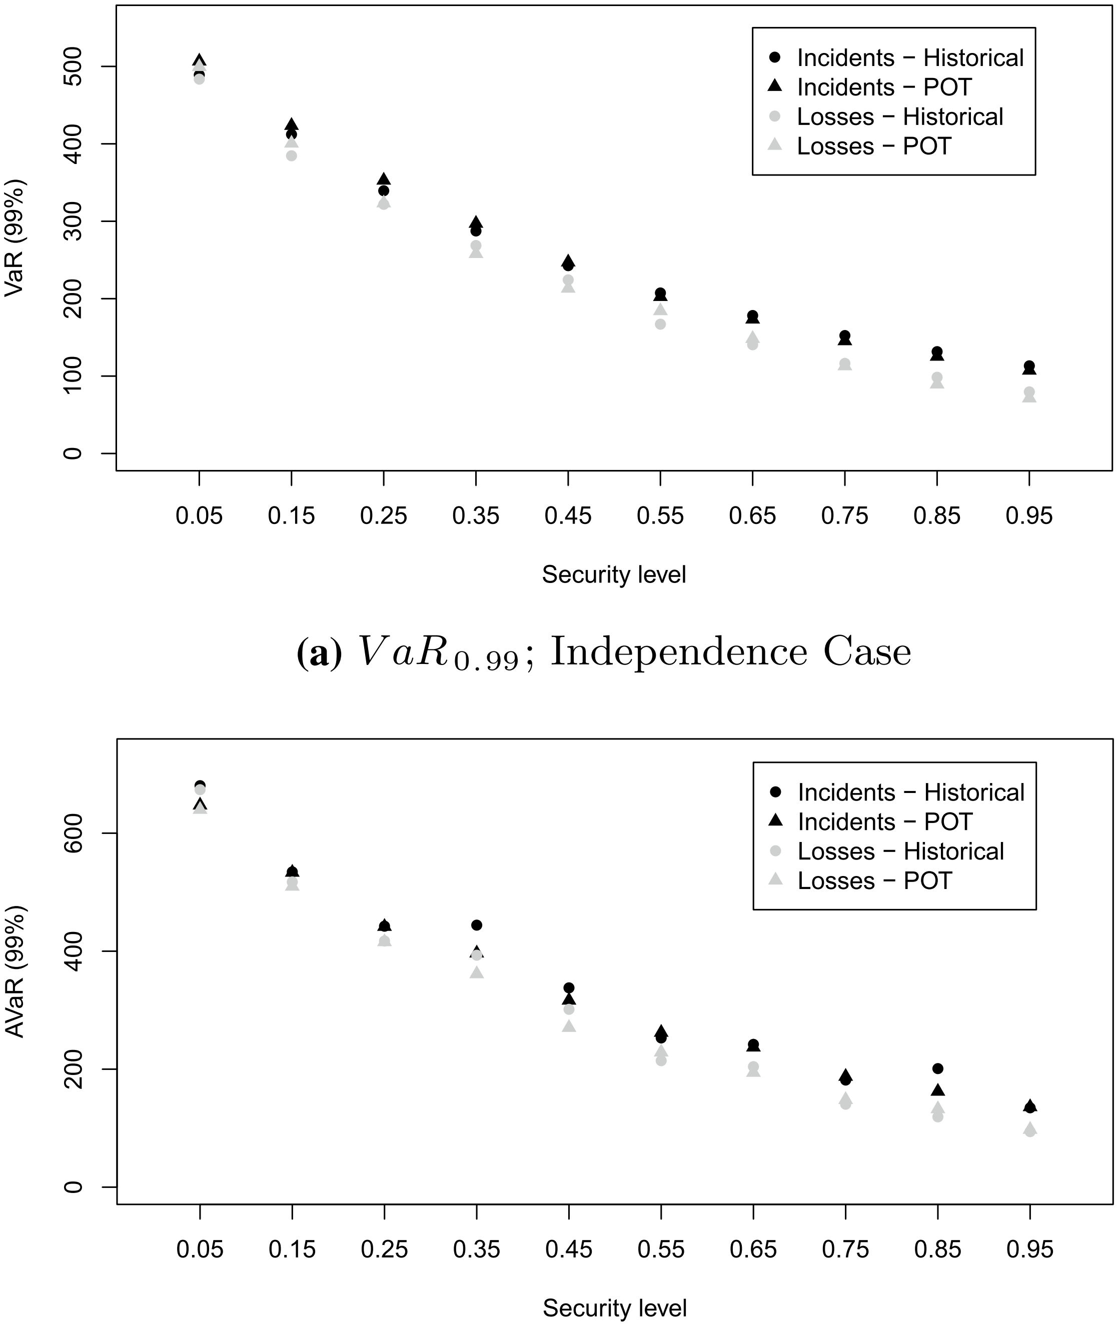

(a) $VaR_{0.99}$; Independence Case

(b) $AVaR_{0.99}$; Independence Case

Fig. 6 Comparison of $VaR_{0.99}$ and $AVaR_{0.99}$ for sub-portfolios of size $K = 50$ with varying security levels

We now assume the cover limit for all contracts to be $\bar{M}_2$ and run the same simulation as before. This could be generalized to allowing different limits depending on the insured's characteristics, e.g. a certain IT security level could be considered a prerequisite for a contract with a high limit. Similarly as above, Fig. 7 displays the (simulated) premium. While changes in the absolute numbers for the expected value principle are minor, the use of other common principles are now viable (all moments exist for the truncated losses) and deliver stable results.

Table 9 Conditional exceedance probabilities $\mathbb{P}(L_{ij} > \bar{M} \mid L_{ij} > u_{ij}) \times 10^2$ of three cover limits for large severity incidents

Cover limit	Low risk	Baseline	High risk
$\bar{M}_1 = 500$	0.0977	0.4055	5.9530
$\bar{M}_2 = 1.000$	0.0437	0.1760	2.1016
$\bar{M}_3 = 10.000$	0.0033	0.0129	0.1335

We observe that a cover limit in most cases impacts only very few (large) claims

Figures showing $VaR_{0.99}$ and $AVaR_{0.99}$ analogously to above are given in Online Appendix A.6. As to be expected, the introduction of a cover limit leads to an overall decrease in both risk measures, where the effect is higher for sub-portfolios with lower security (who tend to suffer the most severe losses and are therefore most impacted by a cover limit) and $AVaR_{0.99}$ decreases more than $VaR_{0.99}$ in absolute numbers (see Figure 9 in Online Appendix A.6).

6 Conclusion

We have presented an actuarial approach to modelling cyber risk that is consistent with the characteristics of the underlying risk factors from an economic and information-technological viewpoint. For this purpose, the existing literature on technical, statistical, economic, actuarial, and legal aspects of cyber risk was analysed in detail to identify relevant risk factors and plausible distributional assumptions within an actuarial framework. By construction, the resulting model is able to capture accumulation risk stemming from multiple firms being simultaneously affected by a cyber event; a prospect that insurers are especially worried about. Some distributional properties of the model and their relevance in the cyber context were highlighted. Moreover, we demonstrated how the model can be implemented in an insurance context using a loss distribution approach. An illustrative simulation study makes use of this implementation and derives the yearly premium for individual contracts as well as common portfolio risk measures. The model is stressed in different directions (contract design, the omission of systemic events) and the findings are analysed from the perspective of an actuary. Given the scarcity of available data on cyber losses, let us reiterate that distributional assumptions and concrete parameter choices rely on the existing literature (scattered across different disciplines) and expert judgments, hence, all quantitative findings should be interpreted with some caution in the light of model/parameter risk. Naturally, since the model presented here is not challenged on data, it is limited to its specific assumptions, e.g. using a Poisson process for arrivals; for the exemplary simulation study, these assumptions are further simplified to illustrate the actuarial exercise. However, to account for updates in the future, we consciously use a modular design that could allow to alter/replace parts of the model or to adapt it to a specific portfolio an insurance company works with.

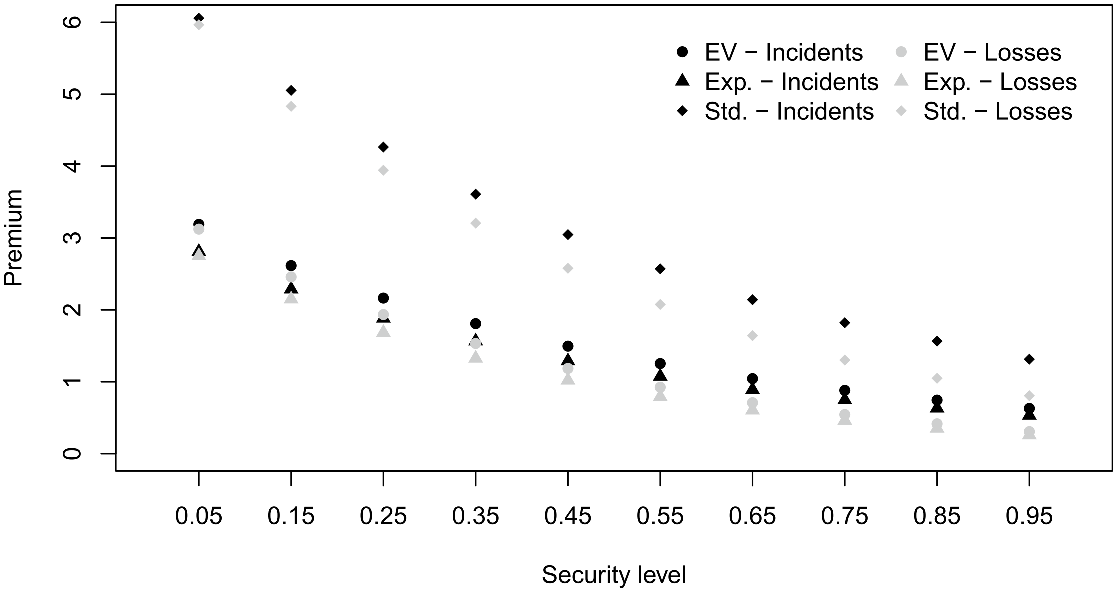

Fig. 7 We compare the premium (according to the three principles introduced in Sect. 4.4) that would be assigned to firms if they were grouped according to their IT security level and every claim in each contract had a limit of $\bar{M}_2$

Many interesting aspects, however, remain open for future research. Once sufficient cyber risk data is available, optimal estimation procedures and out-of-sample tests for the model assumptions are called for. Less theoretical, but equally important, appears the economic/legal question of categorizing cyber incidents. From the actuarial perspective, extremely interesting is the question of (optimal) cyber insurance contract design. Currently offered cyber insurance products seem to reflect the lack of an established common understanding of cyber risk and the resulting caution with which many insurers approach the topic. A better understanding of the underlying dynamics of cyber risk will in time hopefully enable product design to reflect economic optimality criteria instead of the insurers' operational limitations. Furthermore, what separates cyber from most other loss categories is the potential of designing cyber insurance products that transcend mere risk transfer, e.g. by including incident response teams or other services. To the best of our knowledge, this (non-traditional) part of cyber insurance contract design has not yet been addressed from an academic actuarial science viewpoint.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s13385-021-00290-1.

Acknowledgements We have interviewed experts from different fields to gather a comprehensive understanding of cyber risk and insurance. We would especially like to thank the *Cyber-Allianz-Zentrum Bayern (CAZ) am Bayerischen Landesamt für Verfassungsschutz* (Center for Cyber Defense at the Bavarian Office for the Protection of the Constitution) and Frank Romeike for sharing their insights on cyber attacks; as well as Andreas Horn from *ERGO Insurance Group (Cyber-Insurance)* for sharing his knowledge of the current status quo of cyber insurance in Germany and the experts at *Munich Re (Corporate Underwriting Cyber)* for giving their knowledgeable feedback about the practical aspects of the model. We furthermore thank the three anonymous referees whose feedback greatly helped to improve the quality and presentation of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

1. Eling M (2020) Cyber risk research in business and actuarial science. Eur Actuar J 10(2):303–333
2. Anchen J (2017) Cyber: getting to grips with a complex risk. sigma No 1/2017, Swiss Re Institute, Zurich
3. Accenture and Ponemon Institute LLC (2019) The cost of cybercrime: ninth annual cost of cybercrime study unlocking the value of improved cybersecurity protection. https://www.accenture.com/us-en/insights/security/cost-cybercrime-study
4. Allianz Global Corporate & Specialty (2015) A guide to cyber risk: managing the impact of increasing interconnectivity. https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/AGCS-Cyberrisk-report.pdf
5. Ponemon Institute LLC (2016) 2016 Cost of data breach study: global analysis. https://www.academia.edu/35179110/2016_Cost_of_Data_Breach_Study_Global_Analysis
6. Lewis J (2018) Economic impact of cybercrime – no slowing down. McAfee. https://www.mcafee.com/enterprise/en-us/assets/reports/restricted/rp-economic-impact-cybercrime.pdf
7. Sandle P (2018) Ericsson sorry for software glitch that hits mobile services in Britain and Japan. Reuters, 06.12.2018
8. BBC News (2017) Global ransomware attack causes turmoil. BBC, 28.06.2017
9. The Express Tribune (2007) Shadow brokers threaten to release Windows 10 hacking tools. The Express Tribune, 31.05.2017
10. Ponemon Institute LLC (2013) Managing cyber security as a business risk: cyber insurance in the digital age. https://www.ponemon.org/research/ponemon-library/security/managing-cyber-security-as-a-business-risk-cyber-insurance-in-the-digital-age.html
11. Swiss Re and IBM Institute for business value (2016) Cyber: in search of resilience in an interconnected world, Swiss Re Ltd. https://www.swissre.com/dam/jcr:2acf5235-17e1-4a3e-ac71-fec0be9057bf/ZRH-16-09789-P1_Cyber+Publication_web.pdf
12. Allianz Global Corporate & Specialty (2019) Allianz risk barometer - top business risks for 2019. https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/Allianz-Risk-Barometer-2019.pdf
13. ENISA, Robinson N, RAND Europe (2012) Incentives and barriers of the cyber insurance market in Europe. https://www.enisa.europa.eu/publications/incentives-and-barriers-of-the-cyber-insurance-market-in-europe
14. Advisen and PartnerRe (2018) 2018 Survey of cyber insurance market trends. https://partnerre.com/wp-content/uploads/2018/10/2018-Survey-of-Cyber-Insurance-Market-Trends.pdf
15. Romanosky S, Ablon L, Kuehn A, Jones T (2019) Content analysis of cyber insurance policies: how do carriers price cyber risk? J Cybersecur 5(1):117
16. Böhme R, Schwartz G (2010) Modeling cyber-insurance: towards a unifying framework. WEIS, https://www.econinfosec.org/archive/weis2010/papers/session5/weis2010_boehme.pdf
17. Bolot J, Lelarge M (2008) A new perspective on internet security using insurance. In: IEEE INFOCOM 2008 - the 27th conference on computer communications, pp 1948–1956. IEEE, 13.04.2008–18.04.2008

18. Hofmann A (2007) Internalizing externalities of loss prevention through insurance monopoly: an analysis of interdependent risks. Geneva Risk Insur Rev 32(1):91–111
19. Lelarge M, Bolot J (2009) Economic incentives to increase security in the internet: the case for insurance. In: IEEE INFOCOM 2009, pp 1494–1502. IEEE
20. Ogut H, Menon N, Raghunathan S (2005) Cyber insurance and it security investment: impact of interdependence risk. In: WEIS
21. Radosavac S, Kempf J, Kozat U (2008) Using insurance to increase internet security. In: Proceedings of the 3rd international workshop on economics of networked systems, pp 43–48
22. Shetty N, Schwartz G, Walrand J (2010) Can competitive insurers improve network security? In: Trust and trustworthy computing, volume 6101 of Lecture Notes in Computer Science, pp 308–322. Springer
23. Böhme R, Kataria G (2006) Models and measures for correlation in cyber-insurance. WEIS, https://www.econinfosec.org/archive/weis2006/docs/16.pdf
24. Böhme R (2005) Cyber-insurance revisited. WEIS, http://infosecon.net/workshop/pdf/15.pdf
25. Marotta A, Martinelli F, Nanni S, Orlando A, Yautsiukhin A (2017) Cyber-insurance survey. Comput Sci Rev 24:35–61
26. Zhao X, Xue L, Whinston A (2009) Managing interdependent information security risks: a study of cyberinsurance, managed security service and risk pooling. In: ICIS 2009 proceedings, p 49
27. Schwartz G, Sastry S (2014) Cyber-insurance framework for large scale interdependent networks. In: Proceedings of the 3rd international conference on high confidence networked systems, pp 145–154
28. Schwartz G, Shetty N, Walrand J (2013) Why cyber-insurance contracts fail to reflect cyber-risks. In: Başar T, Milenkovic O (eds) 51st Annual Allerton conference on communication, control, and computing (Allerton). IEEE, pp 781–787
29. Shetty N, Schwartz G, Felegyhazi M, Walrand J (2010) Competitive cyber-insurance and internet security. In: Economics of information security and privacy, vol 5, pp 229–247. Springer Science+Business Media LLC
30. Shim W (2012) An analysis of information security management strategies in the presence of interdependent security risk. Asia Pac J Inf Syst 22(1):79–101
31. Pal R (2012) Cyber-insurance for cyber-security a solution to the information asymmetry problem
32. Yang Z, Lui J (2014) Security adoption and influence of cyber-insurance markets in heterogeneous networks. Perform Eval 74:1–17
33. Pal R, Golubchik L, Psounis K, Hui P (2013) On a way to improve cyber-insurer profits when a security vendor becomes the cyber-insurer. In: Proceedings of the 12th IFIP, 2013, pp 1–9. IEEE
34. Pal R, Golubchik L, Psounis K, Hui P (2014) Will cyber-insurance improve network security? A market analysis. In: Proceedings/IEEE INFOCOM, 2014, pp 235–243. IEEE
35. Agrafiotis I, Nurse J, Goldsmith M, Creese S, Upton D (2018) A taxonomy of cyber-harms: defining the impacts of cyber-attacks and understanding how they propagate. J Cybersecur 4(1):tyy006
36. Böhme R, Laube S, Riek M (2018) A fundamental approach to cyber risk analysis. Variance 11(2):161–185
37. Bouveret A (2018) Cyber risk for the financial sector: a framework for quantitative assessment, volume WP/18, 143 of IMF working paper. International Monetary Fund, Washington, DC, June
38. Cebula J, Young L (2010) A taxonomy of operational cyber security risks. Software Engineering Institute, Carnegie Mellon University. https://apps.dtic.mil/dtic/tr/fulltext/u2/a537111.pdf
39. Cohen R, Humphries J, Veau S, Francis R (2019) An investigation of cyber loss data and its links to operational risk. J Oper Risk 14(3):1–25
40. Herath H, Herath T (2011) Copula-based actuarial model for pricing cyber-insurance policies. Insur Mark Co 2(1):7–20
41. Mukhopadhyay A, Chatterjee S, Saha D, Mahanti A, Sadhukhan S (2013) Cyber-risk decision models: to insure it or not? Decis Support Syst 56:11–26
42. Peng C, Xu M, Xu S, Hu T (2018) Modeling multivariate cybersecurity risks. J Appl Stat 45(15):2718–2740
43. Peng C, Xu M, Xu S, Hu T (2017) Modeling and predicting extreme cyber attack rates via marked point processes. J Appl Stat 44(14):2534–2563
44. Xu M, Schweitzer K, Bateman R, Xu S (2018) Modeling and predicting cyber hacking breaches. IEEE Trans Inf Forensics Secur 13(11):2856–2871
45. Baldwin A, Gheyas I, Ioannidis C, Pym D, Williams J (2017) Contagion in cyber security attacks. J Oper Res Soc 68(7):780–791

46. Fahrenwaldt M, Weber S, Weske K (2018) Pricing of cyber insurance contracts in a network model. ASTIN Bull 48(3):1175–1218
47. Xu M, Hua L (2019) Cybersecurity insurance: modeling and pricing. N Am Actuar J 23(2):220–249
48. Xu M, Da G, Xu S (2015) Cyber epidemic models with dependences. Internet Math 11(1):62–92
49. Edwards B, Hofmeyr S, Forrest S (2016) Hype and heavy tails: a closer look at data breaches. J Cybersecur 2(1):3–14
50. Eling M, Loperfido N (2017) Data breaches: goodness of fit, pricing, and risk measurement. Insur Math Econ 75:126–136
51. Eling M, Jung K (2018) Copula approaches for modeling cross-sectional dependence of data breach losses. Insur Math Econ 82:167–180
52. Farkas S, Lopez O, Thomas M (2019) Cyber claim analysis through Generalized Pareto Regression Trees with applications to insurance pricing and reserving. https://hal.archives-ouvertes.fr/hal-02118080
53. Maillart T, Sornette D (2010) Heavy-tailed distribution of cyber-risks. Eur Phys J B 75(3):357–364
54. Wheatley S, Maillart T, Sornette D (2016) The extreme risk of personal data breaches and the erosion of privacy. Eur Phys J B 89(1):59
55. Romanosky S (2016) Examining the costs and causes of cyber incidents. J Cybersecur 2(2):121–135
56. Eling M, Wirfs JH (2019) What are the actual costs of cyber risk events? Eur J Oper Res 272(3):1109–1119
57. Chavez-Demoulin V, Embrechts P, Hofert M (2016) An extreme value approach for modeling operational risk losses depending on covariates. J Risk Insur 83(3):735–776
58. Advisen and PartnerRe (2017) 2017 Survey of cyber insurance market trends. https://partnerre.com/wp-content/uploads/2017/10/PartnerRe-2017-Survey-of-Cyber-Insurance-Market-Trends.pdf
59. Munich Re (2020) Cyber insurance: risks and trends 2020
60. Advisen (2015) 2015 Network security & cyber risk management: the fourth annual survey of enterprise-wide cyber risk management practices in Europe. https://www.advisenltd.com/wp-content/uploads/network-security-cyber-risk-management-white-paper-2015-02-06.pdf
61. Allianz Global Corporate & Specialty (2015) Allianz risk barometer top business risks 2015: risk and reputation in the age of disruption. https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/Allianz-Risk-Barometer-2015.pdf
62. Biener C, Eling M, Wirfs JH (2015) Insurability of cyber risk: an empirical analysis. Geneva Pap Risk Insur Issues Pract 40(1):131–158
63. Eling M, Wirfs JH (2016) Cyber risk: too big to insure? Risk transfer options for a mercurial risk class, volume Band 59 of IVW-HSG-Schriftenreihe. Institute of Insurance Economics I.VW-HSG University of St. Gallen, St. Gallen
64. Gesamtverband der Deutschen Versicherungswirtschaft e.V. (2019) Unverbindlicher Muster-Fragebogen zur Risikoerfassung im Rahmen von Cyber-Versicherungen für kleine und mittelständische Unternehmen. (Unverbindliche Bekanntgabe des Gesamtverbandes der Deutschen Versicherungswirtschaft e.V. (GDV) zur fakultativen Verwendung. Abweichende Vereinbarungen sind möglich.)
65. Allianz Global Corporate & Specialty (2020) Cyber insurance
66. Eling M, Schnell W, Sommerrock F (2016) Ten key questions on cyber risk and cyber risk insurance. The Geneva Association. https://www.genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public//cyber-risk-10_key_questions.pdf
67. Andress J (2014) The basics of information security: understanding the fundamentals of infosec in theory and practice. Syngress, Elsevier Science
68. Carfora M, Martinelli F, Mercaldo F, Orlando A (2019) Cyber risk management: an actuarial point of view. J Oper Risk 14(4):77–103
69. CRO Forum (2016) CRO forum concept paper on a proposed categorisation methodology for cyber risk
70. Deutsche Aktuarvereinigung e.V. Ergebnisbericht des Ausschusses Schadenversicherung: Daten und Methoden zur Bewertung von Cyberrisiken
71. Bandyopadhyay T, Mookerjee V, Rao R (2009) Why IT managers don't go for cyber-insurance products. Commun ACM 52(11):68
72. Jacobs J (2014) Analyzing Ponemon cost of data breach. https://datadrivensecurity.info/blog/posts/2014/Dec/ponemon/
73. Boudreaux D, Rao S, Ferguson W (2013) Measuring losses for small business interruption claims. J Appl Risk Manag Insur 1(1):53–75

74. Deleris L, Elkins D, Pate-Cornell E (2004) Analyzing losses from hazard exposure: a conservative probabilistic estimate using supply chain risk simulation. In: 2004 Winter simulation conference, pp 323–330. IEEE
75. Hashemi SJ, Ahmed S, Khan F (2015) Probabilistic modeling of business interruption and reputational losses for process facilities. Process Saf Prog 34(4):373–382
76. Jain VK, Guin J (2009) Modeling business interruption losses for insurance portfolios. https://www.researchgate.net/publication/290333216
77. Zajdenweber D (1996) Extreme values in business interruption insurance. J Risk Insur 63(1):95
78. Bank for International Settlements BIS (2005) Basel committee on banking supervision. International convergence of capital measurement and capital standards: a revised framework: updated November 2005, volume 118 of Basel Committee Publications. BIS, Basel
79. Daley DJ, Vere-Jones D (2003) An introduction to the theory of point processes: volume I: elementary theory and methods, 2nd edn. Springer, New York
80. Daley DJ, Vere-Jones D (2007) An introduction to the theory of point processes: volume II: general theory and structure. Probability and its applications, 2nd edn. Springer, New York
81. Wood S (2017) Generalized additive models: an introduction with R. Chapman & Hall/CRC texts in statistical science. CRC Press/Taylor & Francis Group, Boca Raton
82. MacKay D (2010) Information theory, inference, and learning algorithms, 1st edn. Cambridge University Press, Cambridge (**9th printing edition, imp. 2010**)
83. Kingman JFC (1993) Poisson processes, vol 3. Oxford studies in probability. Clarendon Press, Oxford
84. Blitzstein J, Hwang J (2015) Introduction to probability. Texts in statistical science, 2nd edn. CRC Press and Taylor & Francis group, Boca Raton
85. Wald A (1944) On cumulative sums of random variables. Ann Math Stat 15(3):283–296
86. Mikosch T (2009) Non-life insurance mathematics: an introduction with the Poisson process. Springer-Verlag, Berlin (**Universitext**)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

A Online Appendix to the paper "A Comprehensive Model for Cyber Risk based on Marked Point Processes and its Application to Insurance"

A.1 Background on Point Processes

Proposition 4 (Superposition, ([49], p.16)). *Let $\{N_i\}_{i\in\mathbb{N}}$ be a countable collection of point processes, then their superposition $\bigcup_{i=1}^{\infty} N_i$ also forms a point process. If $N_1, N_2, \dots$ are independent Poisson processes with mean measures $\Lambda_1, \Lambda_2, \dots$, then their superposition will also be a Poisson process with mean measure $\Lambda = \sum_{i=1}^{\infty} \Lambda_i$.*

Proposition 5 (Thinning ([28], p.34)). *Let $N(\cdot)$ be a (simple, inhomogeneous) Poisson process with rate $\lambda(\cdot)$. Let $p(\cdot)$ be a measurable function on $[0,\infty)$ such that $0 \le p(x) \le 1$ holds $\forall x \in [0,\infty)$. Let a new process $\tilde{N}(\cdot)$ be formed by independently looking at each point of a realization $\{t_i\}$ of $N(\cdot)$ and retaining it with probability $p(x_i)$ (thus deleting it with probability $1-p(x_i)$). Then $\tilde{N}(\cdot)$ is a Poisson process with rate $p(x)\lambda(x)$.*

Definition 1 (Marked Point Process ([28], 6.4.I)). *A marked point process (MPP) with locations in $\mathcal{X}$ and marks in $\mathcal{K}$ is a point process $\{x_i, k_i\}$ on $\mathcal{X} \times \mathcal{K}$ with the additional property that the ground process $N_g(\cdot)$, meaning the process of locations $\{x_i\}$ is itself a point process, i.e. for bounded $A \in \mathcal{B}_{\mathcal{X}}$, $N_g(A) = N(A \times \mathcal{K}) < \infty$.*

Proposition 6 ([28], Prop. 6.4.IV). *Let N be a MPP with independent marks. Then the probability structure of N is completely defined by the distribution of N_g and the mark kernel $\{F(k|x) : k \in \mathcal{B}_{\mathcal{K}}, x \in \mathcal{X}\}$ representing the conditional distribution of the mark, given location x.*

Definition 2 (Compound Poisson process). *Let $N := (N(t))_{t\ge 0}$ be a Poisson process with mean measure $\Lambda(t) > 0$. Let $\{Z_i\}_{i\in\mathbb{N}}$ be a sequence of iid. random variables independent of N. Then the process $R := (R(t))_{t\ge0}$ defined as*

$$R(t) := \sum_{i=1}^{N(t)} Z_i,\ t \ge 0,$$

is called a compound Poisson process.

A.2 Characteristics of Compound Poisson Distribution

Theorem 1 (Wald equation ([78])). *Let $\{X_i\}$ be a sequence of real-valued, iid. random variables and let $N(t) \ge 0$ be an integer-valued r.v. independent of the sequence $\{X_i\}$. Suppose $\mathbb{E}[N(\cdot)] < \infty$ and $\mathbb{E}[X_i] < \infty$. Then*

$$\mathbb{E}\left[\sum_{i=1}^{N(t)} X_i\right] = \mathbb{E}[X_1]\mathbb{E}[N(t)].$$

Theorem 2 (Law of total variance ([15], p. 401)). *Let X and Y be random variables on the same probability space and assume $\mathbb{V}ar[Y] < \infty$. Then*

$$\mathbb{V}ar[Y] = \mathbb{E}[\mathbb{V}ar(Y|X)] + \mathbb{V}ar(\mathbb{E}[Y|X]).$$

The last two results imply that if $\{X_i\}$ is a sequence of iid. random variables and $N(t) \ge 0$ an integer-valued random variable independent of the sequence $\{X_i\}$, then it holds

$$\begin{aligned}\mathbb{V}ar\left(\sum_{i=1}^{N(t)} X_i\right) =: \mathbb{V}ar(Y(t)) &= \mathbb{E}[\mathbb{V}ar(Y(t)|N(t)] + \mathbb{V}ar(\mathbb{E}[Y(t)|N(t)]) \\ &= \mathbb{E}[N(t)\mathbb{V}ar(X_1)] + \mathbb{V}ar(N(t)\mathbb{E}[X_1]) \\ &= \mathbb{V}ar(X_1)\mathbb{E}[N(t)] + \mathbb{E}[X_1]^2\mathbb{V}ar(N(t)).\end{aligned}$$

Proposition 7 ([57], Prop.3.3.4). *Consider the independent compound Poisson sums*

$$L_j = \sum_{i=1}^{N_j} X_i^{(j)}, \qquad j = 1, \dots, K,$$

where $N_j \sim Poi(\lambda_j)$ for some $\lambda_j > 0$ and, for every fixed j, $(X_i^{(j)})_{i=1,2,\dots}$ is an iid. sequence of claim sizes. Then the sum

$$\tilde{L} = L_1 + \ldots + L_K$$

is again compound Poisson with representation

$$\tilde{L} \overset{d}{=} \sum_{i=1}^{N_\lambda} Y_i, \qquad N_\lambda \sim Poi\Big(\sum_{j=1}^{K} \lambda_j\Big),$$

and (Y_i) is an iid. sequence, independent of N_λ, with mixture distribution given by

$$F_{Y_1}(x) = \sum_{j=1}^{K} \frac{\lambda_j}{\sum \lambda_j} F_{X_1^{(j)}}(x), \qquad x \in \mathbb{R}.$$

A.3 Calculations and Proofs from Chapter 4

Proof of Proposition 1. Note that based on *(A4)* for generating S_i and generally m_i being distributed according to cdf. F_M, S_i^* can be thought of as generated analogously to S_i by drawing a realisation of m_i first and then letting

$$\mathbb{P}(Z_{ij} = 1 \mid G_i = 0,\ m_i) = \begin{cases} p_{gen} & \text{iid. } \forall j \in \{1, \ldots, K\} \ s.t. \ c_j < m_i, \\ 0 & \text{else} \end{cases}$$

$$\mathbb{P}(Z_{ij} = 1 \mid G_i = 1, B_i = \hat{b},\ m_i) = \begin{cases} p_{sec} & \text{iid. } \forall j \in \{1 + \sum_{\ell=1}^{\hat{b}-1} K_\ell, \ldots, \sum_{\ell=1}^{\hat{b}} K_\ell\} \ s.t. \ c_j < m_i, \\ 0 & \text{else} \end{cases}$$

i.e. one adds $Z_{ij} \equiv 0$ for all $j : c_j \geq m_i$ in each case, by effectively drawing only on the subset of the portfolio of size $K^* = \max\limits_{k \in \{0, \ldots, K\}} c_{[k]} < m_i$ (resp. the subset of one industry sector $\hat{b}$ of size $K_{\hat{b}}^* = \max\limits_{k \in \{0, \ldots, K_{\hat{b}}\}} c_{[k]}^{\hat{b}} < m_i$).

Conditioning on the realisation of $G_i \in \{0,1\},\ B_i \in \{1, \ldots, B\},\ m_i \in [0,1]$ (in particular, for m_i distinguishing the cases of falling in any of the intervals $[c_{[K^*]}, c_{[K^*+1]}]$ resp. $[c_{[k_{\hat{b}}^*]}^{\hat{b}}, c_{[k_{\hat{b}}^*+1]}^{\hat{b}}]$) yields

$$\mathbb{P}(|S_i^*| = k) = \mathbb{P}(|S_i^*| = k \mid G_i = 0)\ \mathbb{P}(G_i = 0) + \sum_{\hat{b}=1}^{B} \mathbb{P}(|S_i^*| = k \mid G_i = 1,\ B_i = \hat{b})\ \mathbb{P}(B_i = \hat{b} \mid G_i = 1)\ \mathbb{P}(G_i = 1)$$

$$= \underbrace{(1 - p_G) \int_0^1 \mathbb{P}(|S_i^*| = k \mid G_i = 0, m_i = m) \mathrm{d}F_M(m)}_{(I)} + \underbrace{p_G \sum_{\hat{b}=1}^{B} p_{\hat{b}} \int_0^1 \mathbb{P}(|S_i^*| = k \mid G_i = 1, B_i = \hat{b}, m_i = m) \mathrm{d}F_M(m)}_{(II)},$$

where

$$\begin{aligned}
(I) &= (1 - p_G) \int_0^1 \binom{K^*}{k} p_{gen}^k (1 - p_{gen})^{K^* - k} \mathrm{d}F_M(m) \\
&= (1 - p_G) \sum_{K^*=0}^{K} \int_0^1 \mathbb{1}_{[c_{[K^*]}, c_{[K^*+1]}]}(m) \binom{K^*}{k} p_{gen}^k (1 - p_{gen})^{K^* - k} \mathrm{d}F_M(m) \\
&= (1 - p_G) \sum_{K^*=0}^{K} \binom{K^*}{k} p_{gen}^k (1 - p_{gen})^{K^* - k} \big(F_M(c_{[K^*+1]}) - F_M(c_{[K^*]})\big),
\end{aligned}$$

and

$$\begin{aligned}
(II) &= p_G \sum_{\hat{b}=1}^{B} p_{\hat{b}} \int_0^1 \binom{K_{\hat{b}}^*}{k} p_{sec}^k (1 - p_{sec})^{K_{\hat{b}}^* - k} \mathrm{d}F_M(m) \\
&= p_G \sum_{\hat{b}=1}^{B} p_{\hat{b}} \sum_{K_{\hat{b}}^*=0}^{K_{\hat{b}}} \int_0^1 \mathbb{1}_{[c_{[K_{\hat{b}}^*]}^{\hat{b}}, c_{[K_{\hat{b}}^*+1]}^{\hat{b}}]}(m) \binom{K_{\hat{b}}^*}{k} p_{sec}^k (1 - p_{sec})^{K_{\hat{b}}^* - k} \mathrm{d}F_M(m) \\
&= p_G \sum_{\hat{b}=1}^{B} p_{\hat{b}} \sum_{K_{\hat{b}}^*=0}^{K_{\hat{b}}} \binom{K_{\hat{b}}^*}{k} p_{sec}^k (1 - p_{sec})^{K_{\hat{b}}^* - k} \big(F_M(c_{[K_{\hat{b}}^*+1]}^{\hat{b}}) - F_M(c_{[K_{\hat{b}}^*]}^{\hat{b}})\big).
\end{aligned}$$

This implies that $\left|S_i^*\right|$ follows a Binomial mixture distribution, i.e. $f_{\left|S_i^*\right||n,p}(k) = Binom(n,p,k)$ with parameters and weights ($2K + B + 1$ cases):

$$(n,p) = \begin{cases} (K^*, p_{gen}) & \text{with weight} \quad (1-p_G)\Big(F_M(c_{[K^*+1]}) - F_M(c_{[K^*]})\Big), \quad K^* \in \{0,\dots,K\}, \\ (k^*_{\hat{b}}, p_{sec}) & \text{with weight} \quad p_G\, p_{\hat{b}}\, \Big(F_M(c^{\hat{b}}_{[k^*_{\hat{b}}+1]}) - F_M(c^{\hat{b}}_{[k^*_{\hat{b}}]})\Big), \quad k^*_{\hat{b}} \in \{0,\dots,k_{\hat{b}}\},\ \hat{b} \in \{1,\dots,B\}. \end{cases}$$

Again, intuitively this means that, depending on G_i, B_i, and m_i, one draws from a set of different size of potentially affected firms to suffer a loss. As on the respective set, the draws are conditionally iid. Bernoulli draws, the number of “successes” of interest is of course Binomially distributed.
Equation (6) in Proposition 1 follows immediately from above using *(A3)*, i.e. $m_i \sim Unif([0,1])$, thus $F_M(c) = c$, $\forall c \in [0,1]$. Likewise, Equation (5) follows immediately from *(A3)* and by considering the case $c_{[K^*]} = c^{\hat{b}}_{[k^*_{\hat{b}}]} = 0$, $\forall K^* \in \{1,\dots,K\}, \forall k^*_{\hat{b}} \in \{1,\dots,k_{\hat{b}}\}$, $\hat{b} \in \{1,\dots,B\}$. □

Corollary 1 (Moments of number of incidents and losses per event)**.**

$$\begin{aligned}
\mathbb{E}\big[|S_i|\,\big] &= (1-p_G)\, K\, p_{gen} + p_G\, p_{sec} \sum_{\ell=1}^{B} p_\ell\, K_\ell, \\
\mathbb{E}\big[|S_1|^2\,\big] &= (1-p_G)\, \big(K^2 p_{gen}^2 + K\, p_{gen}\, (1-p_{gen})\big) + p_G \sum_{\ell=1}^{B} p_\ell\, \big(K_\ell^2\, p_{sec}^2 + K_l\, p_{sec}\, (1-p_{sec})\big), \\
\mathbb{E}\big[|S_i^*|\,\big] &= (1-p_G)\, p_{gen} \sum_{k^*=0}^{K} k^*\, (c_{[k^*+1]} - c_{[k^*]}) + p_G\, p_{sec} \sum_{\ell=1}^{B} \sum_{k^*_\ell=0}^{K_\ell} p_\ell\, k^*_\ell\, (c^\ell_{[k^*_\ell+1]} - c^\ell_{[k^*_\ell]}), \\
\mathbb{E}\big[|S_i^*|^2\,\big] &= \sum_{k^*=0}^{K} (1-p_G)\, (c_{[k^*+1]} - c_{[k^*]})\, \big((k^*)^2 p_{gen}^2 + k^* p_{gen}(1-p_{gen})\big) \\
&\quad + \sum_{\ell=1}^{B} \sum_{k^*_\ell=0}^{K_\ell} p_G\, p_\ell\, (c^\ell_{[k^*_\ell+1]} - c^\ell_{[k^*_\ell]})\, \big((k^*_\ell)^2 p_{sec}^2 + k^*_\ell p_{sec}(1-p_{sec})\big).
\end{aligned}$$

Proof of Corollary 1. By Proposition 1, $|S_i|$ and $\left|S_i^*\right|$ follow a Binomial mixture distribution.
For $X_i \sim Binom(n,p)$, it holds of course that

$$\begin{aligned}
\mathbb{E}[X_i] &= n\, p, \\
\mathbb{E}[X_i^2] &= n\, p\, (1-p) + n^2 p^2, \\
\mathbb{V}ar[X_i] &= n\, p\, (1-p).
\end{aligned}$$

For a general mixture X of r.v. $\{X_i\}$ with weights $\{w_i\}$, means $\{\mu_i\}$, and variances $\{\sigma_i^2\}$, it holds that

$$\begin{aligned}
\mathbb{E}[X] &= \sum_i w_i \mu_i, \\
\mathbb{E}[X^2] &= \sum_i w_i \mathbb{E}[X_i^2], \\
\mathbb{V}ar[X] &= \Big(\sum_i w_i(\mu_i^2 + \sigma_i^2)\Big) - \mu^2.
\end{aligned}$$

The claims follow directly. □

Lemma 1 (Joint incident and loss probability)**.** *The probability for two firms* $j_1, j_2 \in \{1,\dots,K\}$ *(given their covariates) to register an incident / loss simultaneously from an event is given by*
Case 1: $b_{j_1} = b_{j_2}$ *(same industry sector)*

$$\begin{aligned}
\mathbb{P}(j_1, j_2 \in S_i) &= p_{sec}^2\, p_{b_{j_1}}\, p_G + p_{gen}^2\, (1-p_G), \\
\mathbb{P}(j_1, j_2 \in S_i^*) &= \bar{F}_M\big(\max(c_{j_1}, c_{j_2})\big)\, \big(p_{sec}^2\, p_{b_{j_1}}\, p_G + p_{gen}^2\, (1-p_G)\big).
\end{aligned}$$

Case 2: $b_{j_1} \neq b_{j_2}$ *(different industry sector)*

$$\begin{aligned}
\mathbb{P}(j_1, j_2 \in S_i) &= p_{gen}^2\, (1-p_G), \\
\mathbb{P}(j_1, j_2 \in S_i^*) &= \bar{F}_M\big(\max(c_{j_1}, c_{j_2})\big)\, p_{gen}^2\, (1-p_G).
\end{aligned}$$

Proof of Lemma 1. The statements follow immediately by conditioning and using conditional independence:
Case 1: $b_{j_1} = b_{j_2}$

$$\begin{aligned}
\mathbb{P}(j_1, j_2 \in S_i) &= \mathbb{P}(j_1, j_2 \in S_i \mid G_i = 1, B_i = b_{j_1})\, \mathbb{P}(B_i = b_{j_1} \mid G_i = 1)\, \mathbb{P}(G_i = 1) + \mathbb{P}(j_1, j_2 \in S_i \mid G_i = 0)\, \mathbb{P}(G_i = 0) \\
&= p_{sec}^2\, p_{b_{j_1}}\, p_G + p_{gen}^2\, (1 - p_G), \\
\mathbb{P}(j_1, j_2 \in S_i^*) &= \mathbb{P}\big(j_1, j_2 \in S_i^* \mid G_i = 1, m_i > \max(c_{j_1}, c_{j_2})\big)\, \mathbb{P}\big(G_i = 1 \mid m_i > \max(c_{j_1}, c_{j_2})\big)\, \mathbb{P}\big(m_i > \max(c_{j_1}, c_{j_2})\big) \\
&\quad + \mathbb{P}\big(j_1, j_2 \in S_i^* \mid G_i = 0, m_i > \max(c_{j_1}, c_{j_2})\big)\, \mathbb{P}\big(G_i = 0 \mid m_i > \max(c_{j_1}, c_{j_2})\big)\mathbb{P}\big(m_i > \max(c_{j_1}, c_{j_2})\big) \\
&= p_{sec}^2\, p_{b_{j_1}}\, p_G\, \bar{F}_M\big(\max(c_{j_1}, c_{j_2})\big) + p_{gen}^2\, (1 - p_G)\, \bar{F}_M\big(\max(c_{j_1}, c_{j_2})\big) \\
&= \bar{F}_M\big(\max(c_{j_1}, c_{j_2})\big)\, \big(p_{sec}^2\, p_{b_{j_1}}\, p_G + p_{gen}^2\, (1 - p_G)\big).
\end{aligned}$$

Case 2: $b_{j_1} \neq b_{j_2}$

$$\begin{aligned}
&\mathbb{P}(j_1, j_2 \in S_i) = \mathbb{P}(j_1, j_2 \in S_i \mid G_i = 1)\, \mathbb{P}(G_i = 1) + \mathbb{P}(j_1, j_2 \in S_i \mid G_i = 0)\, \mathbb{P}(G_i = 0) = p_{gen}^2\, (1 - p_G), \\
&\mathbb{P}(j_1, j_2 \in S_i^*) = \bar{F}_M\, \big(\max(c_{j_1}, c_{j_2})\big)\, p_{gen}^2\, (1 - p_G).
\end{aligned}$$

□

Proof of Proposition 2. It follows immediately using Lemma 1

$$\mathbb{P}(j_1 \in S_i \mid j_2 \in S_i) = \frac{\mathbb{P}(j_1, j_2 \in S_i)}{\mathbb{P}(j_2 \in S_i)} = \begin{cases} \frac{p_{sec}^2\, p_{b_{j_2}}\, p_G + p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} = b_{j_2}, \\ \frac{p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})}, & b_{j_1} \neq b_{j_2}, \end{cases}$$

$$\mathbb{P}(j_1 \in S_i^* \mid j_2 \in S_i^*) = \frac{\mathbb{P}(j_1, j_2 \in S_i^*)}{\mathbb{P}(j_2 \in S_i^*)} = \begin{cases} \frac{p_{sec}^2\, p_{b_{j_2}}\, p_G + p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})} & b_{j_1} = b_{j_2},\ c_{j_1} \leq c_{j_2}, \\ \frac{\bar{F}_M(c_{j_1})}{\bar{F}_M(c_{j_2})}\Big(\frac{p_{sec}^2 \cdot p_{b_{j_2}}\, p_G + p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})}\Big) & b_{j_1} = b_{j_2},\ c_{j_1} > c_{j_2}, \\ \frac{p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})} & b_{j_1} \neq b_{j_2},\ c_{j_1} \leq c_{j_2}, \\ \frac{\bar{F}_M(c_{j_1})}{\bar{F}_M(c_{j_2})}\Big(\frac{p_{gen}^2\, (1-p_G)}{\tilde{p}(b_{j_2})}\Big) & b_{j_1} \neq b_{j_2},\ c_{j_1} > c_{j_2}. \end{cases}$$

□

Proof of remark about conditional vs. unconditional probabilities. We have remarked that for firms of the same industry sector, the knowledge about an incident for a firm in the same sector always has a non-negative effect on the incident probability, i.e. for $j_1, j_2 \in \{1, \ldots, K\}$ with $b_{j_1} = b_{j_2} =: b_j$

$$\begin{aligned}
&\mathbb{P}(j_1 \in S_i \mid j_2 \in S_i) \geq \mathbb{P}(j_1 \in S_i \mid b_{j_1}) \\
\overset{\text{Prop.2}}{\iff}\ & p_{sec}^2\, p_{b_j}\, p_G + p_{gen}^2\, (1 - p_G) \geq \big(\tilde{p}(b_j)\big)^2 \\
\iff\ & p_{sec}^2\, p_{b_j}\, p_G + p_{gen}^2\, (1 - p_G) \geq \big(p_G\, p_{b_j}\, p_{sec} + (1 - p_G)\, p_{gen}\big)^2. \qquad (14)
\end{aligned}$$

Generally, for any $\mathbf{x} = (x_1, \ldots, x_n)', \mathbf{y} = (y_1, \ldots, y_n)' \in \mathbf{R}^n$ ($n \in \mathbb{N}$), the Cauchy–Schwarz inequality states that

$$\Big(\sum_{i=1}^n x_i\, y_i\Big)^2 \leq \Big(\sum_{i=1}^n x_i^2\Big)\Big(\sum_{i=1}^n y_i^2\Big).$$

Let $\mathbf{a} = (a_1, \ldots, a_n)' \in (0, \infty)^n, \mathbf{b} = (b_1, \ldots, b_n)' \in \mathbf{R}^n$ ($n \in \mathbb{N}$), and assume $\sum_{i=1}^n a_i \leq 1$. Substituting above $\mathbf{x} = \sqrt{\mathbf{a}}\, \mathbf{b}, \mathbf{y} = \sqrt{\mathbf{a}}$ yields

$$\Big(\sum_{i=1}^n a_i\, b_i\Big)^2 \leq \Big(\sum_{i=1}^n a_i\, b_i^2\Big) \underbrace{\Big(\sum_{i=1}^n a_i\Big)}_{\leq 1} \leq \sum_{i=1}^n a_i\, b_i^2.$$

Substituting for $n = 2$

$$\begin{aligned}
\mathbf{a} &= (a_1, a_2)' = (p_G\, p_{b_j}, (1 - p_G))', \\
\mathbf{b} &= (b_1, b_2)' = (p_{sec}, p_{gen})',
\end{aligned}$$

yields (14). □

Lemma 2 (Moments of cumulative incident and loss numbers)**.** *It holds that*

$$\begin{aligned}
\mathbb{E}\big[\bar{N}^{\cdot,syst}(T)\big] &= \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ K\ p_{gen} + p_G\ p_{sec}\ \sum_{\ell=1}^{B} p_\ell\ K_\ell\Big),\\
\mathbb{V}ar\big[\bar{N}^{\cdot,syst}(T)\big] &= \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ \big(K^2 p_{gen}^2 + K\ p_{gen}\ (1-p_{gen})\big) + p_G \sum_{\ell=1}^{B} p_\ell\ \big(K_\ell^2\ p_{sec}^2 + K_\ell\ p_{sec}\ (1-p_{sec})\big)\Big),\\
\mathbb{E}\big[N^{\cdot,syst}(T)\big] &= \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ p_{gen} \sum_{k^*=0}^{K} k^*\ (c_{[k^*+1]} - c_{[k^*]}) + p_G\ p_{sec} \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_\ell\ k_\ell^*\ (c^\ell_{[k_\ell^*+1]} - c^\ell_{[k_\ell^*]})\Big),\\
\mathbb{V}ar\big[N^{\cdot,syst}(T)\big] &= \Lambda^{\cdot,g}(T)\ \Big(\sum_{k^*=0}^{K} (1-p_G)\ (c_{[k^*+1]} - c_{[k^*]})\ ((k^*)^2 p_{gen}^2 + k^* p_{gen}(1-p_{gen}))\\
&\qquad + \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_G\ p_\ell\ (c^\ell_{[k_\ell^*+1]} - c^\ell_{[k_\ell^*]})\ ((k^*)^2 p_{sec}^2 + k^* p_{sec}(1-p_{sec}))\Big).
\end{aligned}$$

Proof of Lemma 2. For the number of arrivals of the ground process on any interval $[0,T]$, it holds that

$$\mathbb{E}\big[N^{\cdot,g}(T)\big] = \mathbb{V}ar\big[N^{\cdot,g}(T)\big] = \Lambda^{\cdot,g}(T) = \int_0^T \lambda^{\cdot,g}(t)\mathrm{d}t.$$

By Wald's equation and the law of total variance (see Appendix A.2), it follows from Corollary 1

$$\begin{aligned}
\mathbb{E}\big[\bar{N}^{\cdot,syst}(T)\big] &= \mathbb{E}\big[N^{\cdot,g}(T)\big]\ \mathbb{E}\big[|S_i|\big] = \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ K\ p_{gen} + p_G\ p_{sec} \sum_{\ell=1}^{B} p_\ell\ K_\ell\Big),\\
\mathbb{V}ar\big[\bar{N}^{\cdot,syst}(T)\big] &= \mathbb{E}\big[N^{\cdot,g}(T)\big]\ \mathbb{V}ar\big[|S_i|\big] + \mathbb{E}\big[|S_i|\big]^2\ \mathbb{V}ar[N^{\cdot,g}(T)] = \mathbb{E}\big[N^{\cdot,g}(T)\big]\ \mathbb{E}\big[|S_i|^2\big]\\
&= \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ \big(K^2 p_{gen}^2 + K\ p_{gen}\ (1-p_{gen})\big) + p_G \sum_{\ell=1}^{B} p_\ell\ \big(K_\ell^2\ p_{sec}^2 + K_\ell\ p_{sec}\ (1-p_{sec})\big)\Big).
\end{aligned}$$

Likewise,

$$\begin{aligned}
\mathbb{E}\big[N^{\cdot,syst}(T)\big] &= \mathbb{E}\big[N^{\cdot,g}(T)\big]\ \mathbb{E}\big[|S_i^*|\big]\\
&= \Lambda^{\cdot,g}(T)\ \Big((1-p_G)\ p_{gen} \sum_{k^*=0}^{K} k^*\ (c_{[k^*+1]} - c_{[k^*]}) + p_G\ p_{sec} \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_\ell\ k_\ell^*\ (c^\ell_{[k_\ell^*+1]} - c^\ell_{[k_\ell^*]})\Big),
\end{aligned}$$

$$\begin{aligned}
\mathbb{V}ar\big[N^{\cdot,syst}(T)\big] &= \mathbb{E}\big[N^{\cdot,g}(T)\big]\ \mathbb{E}\big[|S_1^*|^2\big]\\
&= \Lambda^{\cdot,g}(T)\ \Big(\sum_{k^*=0}^{K} (1-p_G)\ (c_{[k^*+1]} - c_{[k^*]})\ ((k^*)^2 p_{gen}^2 + k^* p_{gen}(1-p_{gen}))\\
&\quad + \sum_{\ell=1}^{B} \sum_{k_\ell^*=0}^{K_\ell} p_G\ p_\ell\ (c^\ell_{[k_\ell^*+1]} - c^\ell_{[k_\ell^*]})\ ((k^*)^2 p_{sec}^2 + k^* p_{sec}(1-p_{sec}))\Big).
\end{aligned}$$

□

Proof of Proposition 3. It follows immediately from Lemma 2 that

$$\begin{aligned}
DI(\bar{N}^{\cdot,syst}(T)) &= \frac{\mathbb{V}ar\big[\bar{N}^{\cdot,syst}(T)\big]}{\mathbb{E}\big[\bar{N}^{\cdot,syst}(T)\big]} = \frac{\mathbb{E}\big[|S_i|^2\big]}{\mathbb{E}\big[|S_i|\big]}\\
&= \frac{(1-p_G)\ \big(K^2 p_{gen}^2 + K\ p_{gen}\ (1-p_{gen})\big) + p_G \sum_{\ell=1}^{B} p_\ell\ (K_\ell^2\ p_{sec}^2 + K_\ell\ p_{sec}\ (1-p_{sec}))}{(1-p_G)\ K\ p_{gen} + p_G\ p_{sec} \sum_{\ell=1}^{B} p_\ell\ K_\ell}\\
&= 1 + \frac{(1-p_G)\ p_{gen}^2\ (K^2 - K) + p_G\ p_{sec}^2 \sum_{\ell=1}^{B} p_\ell (K_\ell^2 - K_\ell)}{(1-p_G)\ K\ p_{gen} + p_G\ p_{sec} \sum_{\ell=1}^{B} p_\ell\ K_\ell} > 1,
\end{aligned}$$

and likewise

$$
\begin{aligned}
DI\big(N^{\cdot,syst}(T)\big) &= \frac{\mathbb{V}ar\big[N^{\cdot,syst}(T)\big]}{\mathbb{E}\big[N^{\cdot,syst}(T)\big]} = \frac{\mathbb{E}\big[\big|S_i^*\big|^2\big]}{\mathbb{E}\big[\big|S_i^*\big|\big]} \\
&= \frac{\sum_{k^*=0}^{K}(1-p_G)\ (c_{[k^*+1]}-c_{[k^*]})\ ((k^*)^2 p_{gen}^2 + k^* p_{gen}(1-p_{gen}))}{(1-p_G)\ p_{gen}\sum_{k^*=0}^{K} k^*\ (c_{[k^*+1]}-c_{[k^*]}) + p_G\ p_{sec}\sum_{\ell=1}^{B}\sum_{k_\ell^*=0}^{K_\ell} p_\ell\ k_\ell^*\ (c^\ell_{[k_\ell^*+1]}-c^\ell_{[k_\ell^*]})} \\
&\quad + \frac{\sum_{\ell=1}^{B}\sum_{k_\ell^*=0}^{K_\ell} p_G\ p_\ell\ (c^\ell_{[k_\ell^*+1]}-c^\ell_{[k_\ell^*]})\ ((k_\ell^*)^2 p_{sec}^2 + k_\ell^* p_{sec}(1-p_{sec}))}{(1-p_G)\ p_{gen}\sum_{k^*=0}^{K} k^*\ (c_{[k^*+1]}-c_{[k^*]}) + p_G\ p_{sec}\sum_{\ell=1}^{B}\sum_{k_\ell^*=0}^{K_\ell} p_\ell\ k_\ell^*\ (c^\ell_{[k_\ell^*+1]}-c^\ell_{[k_\ell^*]})} \\
&= 1 + \frac{(1-p_G)\sum_{k^*=0}^{K} p_{gen}^2((k^*)^2-k^*)(c_{[k^*+1]}-c_{[k^*]}) + p_G\sum_{\ell=1}^{B}\sum_{k_\ell^*=0}^{K_\ell} p_{sec}^2 p_\ell((k_\ell^*)^2-k_\ell^*)(c^\ell_{[k_\ell^*+1]}-c^\ell_{[k_\ell^*]})}{(1-p_G)p_{gen}\sum_{k^*=0}^{K} k^*(c_{[k^*+1]}-c_{[k^*]}) + p_G p_{sec}\sum_{\ell=1}^{B}\sum_{k_\ell^*=0}^{K_\ell} p_\ell k_\ell^*(c^\ell_{[k_\ell^*+1]}-c^\ell_{[k_\ell^*]})} > 1.
\end{aligned}
$$

The fractions in the last lines of both equations are obviously non-negative (positive, under the additional assumptions in Proposition 3), as they contain only sums and products of non-negative (positive) quantities. □

A.4 Alternative Severity Distributions

DB: Use link between number of records and cost

While it is difficult to find reliable empirical data about the cost of a data breach, some data about the number of breached / stolen records is available. Thus, several authors have tried to find a link between the number of records and the cost of a data breach. The often cited *Jacob's formula* ([47]) suggests to link the log-transformed cost L of a data breach to the number of compromised records D according to

$$\log(L) = 7.68 + 0.7584\ \log(D). \tag{15}$$

An amendment to this formula was proposed in [41], who argue that [47] did not yet take into account the cost of *mega data breaches* observed in future years and thus alternatively propose

$$\log(L) = -1.998 + 7.503\ \log\big(\log(D)\big). \tag{16}$$

Therefore, an alternative to modelling the cost of a data breach directly using a combination of log-normal and GPD would be to first model the number of breached records using a log-normal (as suggested by the results in [32]) and then convert the number of records into monetary losses using (15) or (16).

In the context of this work, let D_{ij} be the number of lost / stolen records in a DB incident at time t_i affecting firm j (where $\{t_i\}_{i\in\mathbb{N}}$ only counts the event times at firm j). Then assume

$$
\begin{aligned}
D_{ij} &\sim LN\big(\mu_j^{DB}(t_i), \sigma_j^{DB}(t_i)\big), \qquad (17)\\
\mu_j^{DB}(t_i) &= \alpha_{\mu,DB} + f_{\mu,DB,3}(x_{j3}) + f_{\mu,DB,4}(x_{j4}) + g_{\mu,DB}(t_i),\\
\sigma_j^{DB}(t_i) &\equiv \sigma^{DB},
\end{aligned}
$$

where the functions $f_{\mu,DB,\cdot}$ and $g_{\mu,DB}$ are as usual. By (15), the number of records D_{ij} is converted into the cost of the breach L_{ij} according to

$$\log(L_{ij}) = 7.68 + 0.7584\ \log(D_{ij}),$$

which is equivalent to directly assuming that

$$
\begin{aligned}
L_{ij} &\sim LN(\hat{\mu}_j^{DB}(t_i), \hat{\sigma}^{DB}),\\
\hat{\mu}_j^{DB}(t_i) &= \alpha_{\hat{\mu},DB} + f_{\hat{\mu},DB,3}(x_{j3}) + f_{\hat{\mu},DB,4}(x_{j4}) + g_{\hat{\mu},DB}(t_i),\\
\hat{\sigma}_j^{DB}(t_i) &\equiv \hat{\sigma}^{DB}
\end{aligned}
$$

Likewise, using (16) to convert the number of records into the cost of the breach, i.e. assume D_{ij} to be distributed according to (17) and the breach cost L_{ij} then to be given by

$$\log(L_{ij}) = -1.998 + 7.503 \log\big(\log(D_{ij})\big).$$

BI: Replace log-normal by PERT

Regarding the economic impact of BI incidents, some sources from the non-cyber domain are available ([21, 30, 44, 48, 85]). The only sources including indications of which distributions are useful to model economic loss from BI are [85], who finds the size of yearly BI insurance claims to follow a Pareto distribution with an extremely heavy tail and infinite expected claim size, and [44], who suggests modelling BI loss by a *PERT* distribution, a special case of the beta distribution with the three parameters minimum x_{min}, mode x_{mode}, and maximum x_{max} with density

$$\begin{aligned} f_{PERT}(x) &= \frac{(x-x_{min})^{v-1}(x_{max}-x)^{w-1}}{Beta(v,w)(x_{max}-x_{min})^{v+w-1}} 1_{[x_{min},x_{max}]}, \\ v &= 1+\gamma_P\Big(\frac{x_{mode}-x_{min}}{x_{max}-x_{min}}\Big), \\ w &= 1+\gamma_P\Big(\frac{x_{max}-x_{mode}}{x_{max}-x_{min}}\Big), \end{aligned}$$

where $Beta(\cdot)$ is the Beta function and for the standard PERT $\gamma_P = 4$.
Thus, for BI incidents, one could suggest replacing the log-normal distribution for the body by a PERT distribution, i.e. assume for a BI loss L_{ij} at time t_i affecting firm j it holds

$$\begin{aligned} (L_{ij} \mid L_{ij} \leq u_{ij}^{BI}) &\sim PERT(x_{ij}^{\min}, x_{ij}^{mode}, x_{ij}^{\max}, 4), \\ x_{ij}^{\min} &= 0, \\ x_{ij}^{\max} &= u_{ij}^{BI}, \\ x_{ij}^{\text{mode}} &= \exp\big(\mu_j^{BI}(t_i) - \sigma_j^{BI}(t_i)^2\big), \end{aligned}$$

where $PERT(x^{\min}, x^{\text{mode}}, x^{\max}, 4)$ denotes the *PERT* distribution with minimum, mode, and maximum values $x^{\min}, x^{\text{mode}}, x^{\max}$ respectively and standard shape parameter $\gamma_p = 4$. The mode and threshold (maximum) are chosen such that they coincide with the ones from the underlying log-normal used to find the threshold between body and tail of the loss distribution.

A.5 Comparison of all Premium Calculation Results

Below, we compare the premiums (for three individual firms and all sub-portfolios) obtained from the simulation with dependent losses, with independent losses, with cover limit, and with the premiums obtained from calculating the (discretized) loss distribution pdf. using Panjer recursion. We observe that they are very similar in all cases; as for the latter two cases (simulation with cover limit and Panjer recursion) loss severities are truncated from above, the application of premium principles that depend on more than just the first moment are feasible and the results for the expected value principle are slightly lower.

Premium Principle			
	Expected Value ($\rho = 0.2$)	Exponential ($\gamma = 10^{-3}$)	Standard Deviation ($\rho = 0.2$)
Premium based on dependent simulated losses (incidents)			
Firm 1	2.0814 (2.2338)	–	–
Firm 2	0.4451 (0.7746)	–	–
Firm 3	1.1732 (1.5164)	–	–
Premium based on independent simulated losses (incidents)			
Firm 1	2.1213 (2.2514)	–	–
Firm 2	0.4726 (1.3395)	–	–
Firm 3	1.2149 (1.5353)	–	–
Premium based on simulated losses (incidents) with cover limit			
Firm 1	2.1592 (2.3051)	1.8993 (2.0258)	4.5101 (4.7022)
Firm 2	0.4385 (0.7783)	0.3717 (0.6608)	1.0745 (1.6300)
Firm 3	1.1620 (1.5115)	0.9960 (1.2956)	2.4413 (2.9404)
Premium based on Panjer recursion			
Firm 1	1.7633 (1.8849)	2.1160 (2.2619)	1.8366 (1.9632)
Firm 2	0.3797 (0.662)	0.4557 (0.7944)	0.3861 (0.6731)
Firm 3	0.9605 (1.2643)	1.1526 (1.5171)	0.9874 (1.2997)

Table 10: Comparison of one-year cyber insurance premiums for three selected firms, based on 50.000 simulation runs (upper panels) and Panjer recursion for the given assumptions and parameter values (lower panel). Numbers in brackets indicate what the premium would have been if all incoming incidents had been counted.

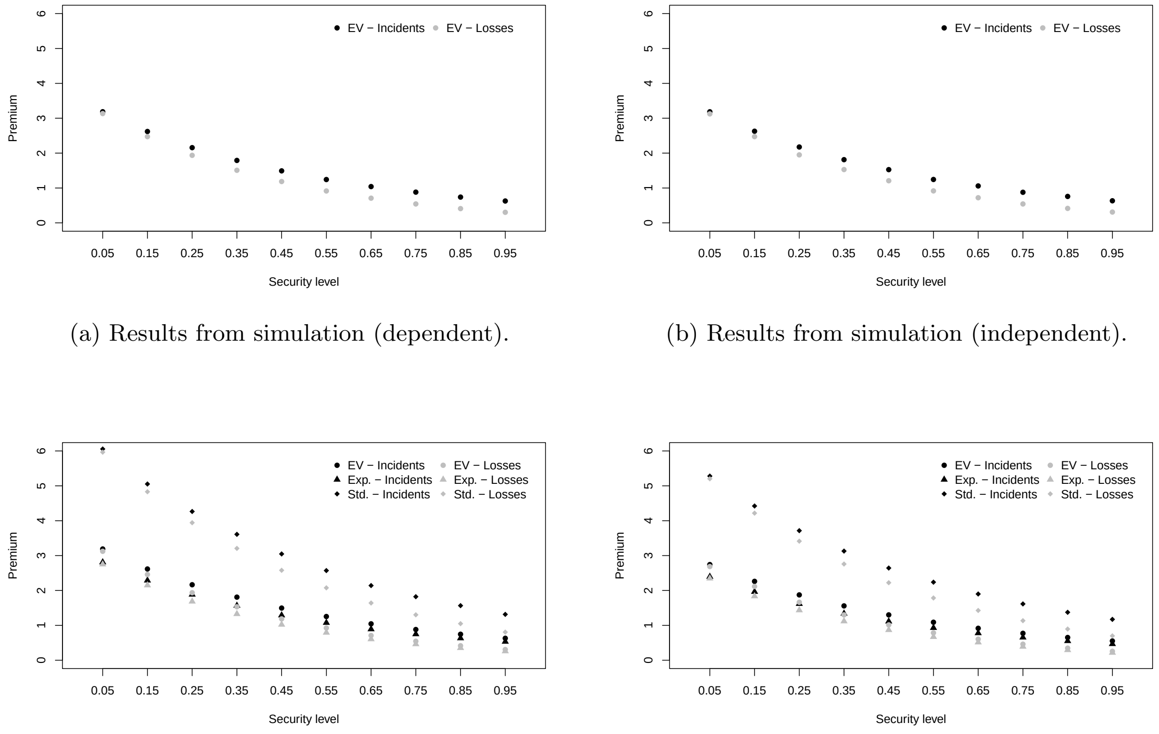

(a) Results from simulation (dependent).

(b) Results from simulation (independent).

(c) Results from simulation (with cover limit).

(d) Results from Panjer recursion.

Figure 8: We compare the premium that would be assigned to firms if they were grouped according to their IT security level, based on the three simulation studies and the implemented Panjer recursion scheme. Note that the results in Figure 8d are for firms with the given security level and otherwise baseline covariate levels, so should be expected slightly below results from simulations of the sub-portfolios with mixed covariate levels.

A.6 Risk Measures for Simulation with Cover Limit

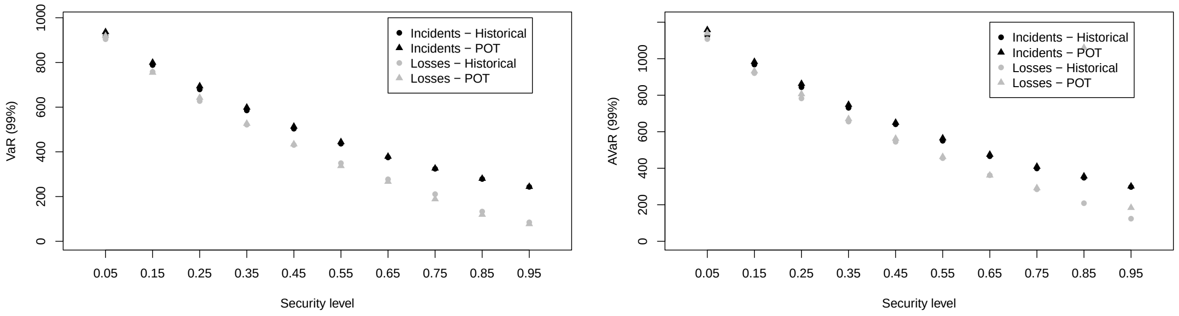

(a) $VaR_{0.99}$; Cover Limit $\bar{M}_2$.

(b) $AVaR_{0.99}$; Cover Limit $\bar{M}_2$.

Figure 9: Comparison of $VaR_{0.99}$ and $AVaR_{0.99}$ for all sub-portfolios.

A.7 Details on Covariate Levels

	Number of employees		
Revenue	small	medium	large
small	1	1	(2)
medium	1	2	3
large	(2)	3	3

Table 11: Factor levels of s by combinations of revenue and number of employees. This is in line with the classification of *SMEs* in the European Union. As revenue and number of employees are highly correlated, only very few companies should fall into the classifications in the upper right or lower left cell.

	Number of stored records	
Sensitive Data	$\leq$ threshold	$>$ threshold
No	1	2
Yes	2	3

Table 12: Factor levels of d, given the number of stored records and sensitivity of data. Sensitive data includes, e.g. *Personally Identifiable Information (PII)*, *Protected Health Information (PHI)*, or classified government data. Despite the labels, this is not a numerical attribute and it is not clear whether the two cases labeled 2 (medium risk) are comparable, or, if considered not comparable, how they should be ordered.

	Number of employees e		
Sector b	small	medium	large
HC, EDU, GOV	1	1	2
FI, BR, MAN	1	2	3

Table 13: Factor levels of *nsup* by combinations of sector and number of employees. The classification relies on expert judgment and is not founded by empirical evidence. An insurance company might simply obtain this information from its customers.

A.2 Risk mitigation services in cyber insurance: optimal contract design and price structure [2]

Summary

This article investigates novel cyber insurance policies equipped with risk mitigation services, in particular how to optimally price such contracts from an insurer's viewpoint. This is a highly relevant topic on the currently evolving cyber insurance market, as policies continue to mature and the potential of including so-called *cyber assistance*, i.e. *pre-incident* and *post-incident services* is recognized by insurers and insurance buyers (e.g. [104]). However, to the best of our knowledge, these policies are currently priced on an ad-hoc basis and established actuarial pricing approaches are yet to be developed. Theoretically, policies combining risk transfer and risk mitigation give rise to interesting trade-offs for both parties: The insurance buyer balances (deterministic) upfront investment into risk mitigation against some reduction of (stochastic) future risk, and the insurer weighs some reduction of the (stochastic) to-be-insured (portfolio) risk against a reduced (deterministic) feasible premium.

The first part of the article outlines types of cyber assistance and connects services offered on the market to the established theoretical concepts of *self-protection* and *self-insurance* ([51]). The former refers to activities modifying the probability of a loss, whereas the latter encompasses activities shaping the loss severity if a loss manifests. We thus identify existing pre-incident and post-incident services as self-protection and self-insurance activities, respectively, and furthermore suggest a new type of self-protection activity requiring a portfolio viewpoint, to use dependencies between risks to all parties' advantage. Next, we introduce the prerequisites of the mathematical model inspired by the framework of [23], i.e. to capture the effect of risk mitigation services on the loss distribution by inducing a decreasing order in the sense of first-order stochastic dominance (see, e.g., [103]), to model the risk measurement of both parties by (concave) distortion risk measures (e.g. [136]) and their interaction as a so-called Stackelberg game (e.g. [76]). Within this sequential optimization game, the insurer *leads* by choosing a price structure, i.e. a combination of risk premium and service premium, which determines the price of any offered contract, and the buyer *follows* by selecting a contract, determined by the level of risk mitigation services and proportional insurance share. This part concludes by formalising both parties' optimization problems and connecting the insurance buyer's choices to classical ways of dealing with risk (e.g. [100]).

A typical approach to a Stackelberg game is by *backward induction* (see e.g. [111]). Thus, the next two sections deal with first deriving the insurance buyer's optimal response to any insurer's choice of price structure, and subsequently finding the insurer's optimal solution (given knowledge of all the buyer's optimal responses). In this part, we still consider a single-contract perspective. The buyer's problem is conceptually similar to [23] (but slightly more involved in the self-insurance case due to the different structure of the loss function), such that we focus on the novelties originating from the new formulation of the insurer's problem and the interpretation of all results in the cyber insurance context. The insurer's problem is now a more general bivariate problem (choosing a combination of risk premium and service premium) and includes the univariate analysis in [23] as a special case. The main finding of this section is that under the given assumptions, the insurer would never subsidize services in a self-protection scenario (i.e. it is always an optimal solution to shift the full cost of risk mitigation services to the buyer), but might indeed do so in a self-insurance scenario. This conceptual difference stems from the converse extent of the effect of services on the price of insurance and the distortion risk measure in the two cases. The results are illustrated in extensive case studies (in the Online Appendix). We extend the analysis by considering the self-protection scenario from a portfolio viewpoint, and show that the results from the single-contract case do not necessarily carry over, i.e. already in simple bivariate examples of dependent risks, it may be optimal for the insurer to subsidize risk mitigation services for some policyholders. We consider dependence mechanisms representative for cyber risk, namely directed loss propagation (originating from network models, e.g. [71, 145]), common events (e.g. [29, 4]) and copula approaches (e.g. [54, 114]). We present an exemplary extension to a general multivariate model for the case of common cyber events (in the Online Appendix).

In summary, this article extends the landscape of previous studies on the combination of risk mitigation and risk transfer by bestowing the insurer with a more central role, namely controlling the price of both activities. This relates to the real-world situation in cyber insurance, where insurers have started to endow insurance policies with cyber assistance services.

The contribution offers threefold insights, regarding the viewpoints of insurers, (prospective) insurance buyers, and the general cyber insurance market. For insurers, the study of the insurer's new bivariate optimization problem offers first guidance to the optimal pricing of insurance policies including cyber assistance. For insurance buyers, it is invaluable to better understand how different contracts would be optimally priced by an insurer. In particular, it is relevant that the price structure which a prospective policyholder is offered (and the implicit incentive for risk mitigation) may not only depend on his own characteristics, but on the insurer's existing portfolio and the dependence between risks. Finally, the study of the insurance buyer's problem may serve to theoretically explain the insurance gap observed in the cyber insurance market (see, e.g., [127]), and suggests equipping insurance policies with (potentially subsidized) services, which help to alter the risk in a way that allows the insurer to offer desired coverage at an acceptable (from the buyer's viewpoint) premium, as a way to mitigate it.

Reception

The work contained in this article (prior to publication) has been presented at several scientific conferences and has received the *First Prize of Young Researcher Best Presentation Awards* at the *11th Conference in Actuarial Science & Finance on Samos* as well as the *FiVeG Award 2022 for the Best Junior Presentation* at the *13th CEQURA Conference on Advances in Financial and Insurance Risk Management.*

Individual contributions

I am the main author of this article. The idea of investigating this topic was developed jointly with my supervisor Matthias Scherer, who also made helpful suggestions regarding both content and presentation of the article during our regular discussions. I was responsible for the writing of the manuscript (the whole first draft as well as subsequent drafts based on comments by Matthias Scherer), developing the proofs of all statements contained therein, and the implementation and presentation of the case studies.

Permission to include the article

Licence to Publish - Open Access

SPRINGER NATURE

Licensee:	The Geneva Association	(the 'Licensee')
Journal Name:	The Geneva Papers on Risk and Insurance - Issues and Practice	(the 'Journal')
Manuscript Number:	GPRI-D-22-00043R2	
Proposed Title of Article:	Risk mitigation services in cyber insurance: optimal contract design and price structure	(the 'Article')
Author(s) [Please list all named Authors]:	Gabriela Zeller, Matthias Scherer	(the 'Author')
Corresponding Author Name:	Gabriela Zeller	

Licence Applicable to the Article:

Creative Commons licence CC BY: This licence allows readers to copy, distribute and transmit the Article as long as it is attributed back to the author. Readers are permitted to alter, transform or build upon the Article, and to use the Article for commercial purposes. Please read the full licence for further details at - http://creativecommons.org/licenses/by/4.0/

Subject to editorial acceptance of the Article, it will be published under the Creative Commons licence shown above.

1 Grant of Rights

a) For good and valuable consideration, the Author hereby grants to the Licensee the perpetual, non-exclusive, irrevocable, world-wide, assignable, sublicensable and unlimited right to: publish, reproduce, copy, distribute, communicate, display publicly, sell, rent and/ or otherwise make available the article identified above, including any supplementary information and graphic elements therein (e.g. illustrations, charts, moving images) (the "Article") in any language, in any versions or editions in any and all forms and/or media of expression (including without limitation in connection with any and all end-user devices), whether now known or developed in the future. Without limitation, the above grant includes: (i) the right to edit, alter, adapt, adjust and prepare derivative works; (ii) all commercial use, advertising, and marketing rights, including without limitation graphic elements on the cover of the journal and in relation to social media; (iii) rights for any training, educational and/or instructional purposes; (iv) the right to add and/or remove links or combinations with other media/works; and (v) the right to create, use and/or license and/or sublicense content data or metadata of any kind in relation to the Article (including abstracts and summaries) without restriction. The above rights are granted in relation to the Article as a whole or any part and with or in relation to any other works.

b) Without limiting the rights granted above, Licensee is granted the rights to use the Article for the purposes of analysis, testing, and development of publishing- and research-related workflows, systems, products, projects, and services; to confidentially share the Article with select third parties to do the same; and to retain and store the Article and any associated correspondence/files/forms to maintain the historical record, and to facilitate research integrity investigations. The grant of rights set forth in this clause (b) is irrevocable.

c) The Licensee will have the right, but not the obligation, to exercise any or all of the rights granted herein. If the Licensee elects not to publish the Article for any reason, all publishing rights under this Agreement as set forth in clause 1.a) above will revert to the Author.

2 Copyright

Ownership of copyright in the Article will be vested in the name of the Author. When reproducing the Article or extracts from it, the Author will acknowledge and reference first publication in the Journal.

3 Use of Article Versions

a) For purposes of this Agreement: (i) references to the "Article" include all versions of the Article; (ii) "Submitted Manuscript" means the version of the Article as first submitted by the Author; (iii) "Accepted Manuscript" means the version of the Article accepted for publication, but prior to copy-editing and typesetting; and (iv) "Version of Record" means the version of the Article published by the Licensee, after copy-editing and typesetting. Rights to all versions of the Manuscript are granted on a non-exclusive basis.

b) The Author may make the Submitted Manuscript available at any time and under any terms (including, but not limited to, under a CC BY licence), at the Author's discretion. Once the Article has been published, the Author will include an acknowledgement and provide a link to the Version of Record on the publisher's website: "This preprint has not undergone peer review (when applicable) or any post-submission improvements or corrections. The Version of Record of this article is published in [insert journal title], and is available online at https://doi.org/[insert DOI]".

c) Immediately after acceptance the Author may deposit the Accepted Manuscript to any location, and under any terms (including, but not limited to, under a CC BY licence), provided it is not made publicly available until after publication. The Author will include an acknowledgement in the Accepted Manuscript, together with a link to the Version of Record on the publisher's website: "This version of the article has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/[insert DOI]".

4 Warranties & Representations

Author warrants and represents that:

a)

i. the Author is the sole copyright owner or has been authorised by any additional copyright owner(s) to grant the rights defined in clause 1,

ii. the Article does not infringe any intellectual property rights (including without limitation copyright, database rights or trade mark rights) or other third party rights and no licence from or payments to a third party are required to publish the Article,

iii. the Article has not been previously published, nor has the Author committed to licensing any version of the Article under a licence inconsistent with the terms of this Agreement,

iv. if the Article contains materials from other sources (e.g. illustrations, tables, text quotations), Author has obtained written permissions to the extent necessary from the copyright holder(s), to license to the Licensee the same rights as set out in clause 1 and has cited any such materials correctly;

b) all of the facts contained in the Article are according to the current body of research true and accurate;

c) nothing in the Article is obscene, defamatory, violates any right of privacy or publicity, infringes any other human, personal or other rights of any person or entity or is otherwise unlawful and that informed consent to publish has been obtained for any research participants;

d) nothing in the Article infringes any duty of confidentiality owed to any third party or violates any contract, express or implied, of the Author;

e) all institutional, governmental, and/or other approvals which may be required in connection with the research reflected in the Article have been obtained and continue in effect;

f) all statements and declarations made by the Author in connection with the Article are true and correct; and

g) the signatory who has signed this agreement has full right, power and authority to enter into this agreement on behalf of all of the Authors.

5 Cooperation

a) The Author will cooperate fully with the Licensee in relation to any legal action that might arise from the publication of the Article, and the Author will give the Licensee access at reasonable times to any relevant accounts, documents and records within the power or control of the Author. The Author agrees that any Licensee affiliate through which the Licensee exercises any rights or performs any obligations under this Agreement is intended to have the benefit of and will have the right to enforce the terms of this Agreement.

b) Author authorises the Licensee to take such steps as it considers necessary at its own expense in the Author's name(s) and on their behalf if the Licensee believes that a third party is infringing or is likely to infringe copyright in the Article including but not limited to initiating legal proceedings.

6 Author List

Changes of authorship, including, but not limited to, changes in the corresponding author or the sequence of authors, are not permitted after acceptance of a manuscript.

7 Post Publication Actions

The Author agrees that the Licensee may remove or retract the Article or publish a correction or other notice in relation to the Article if the Licensee determines that such actions are appropriate from an editorial, research integrity, or legal perspective.

8 Controlling Terms

The terms of this Agreement will supersede any other inconsistent terms that the Author or any third party may assert apply to any version of the Article.

9 Governing Law

This Agreement will be governed by, and construed in accordance with, the laws of England and Wales. The courts of England, UK will have exclusive jurisdiction.

Signed for and on behalf of the Author(s)
Corresponding Author: Gabriela Zeller
Email: gabi.zeller@tum.de
IP Address: 88.128.92.101

Time Stamp: 2023-03-02 09:24:27

Springer SpringerLink shop

Subjects Services About Us

Permissions

Get permission to reuse Springer Nature content

Springer Nature is partnered with the Copyright Clearance Center to meet our customers' licensing and permissions needs.

Copyright Clearance Center’s RightsLink® service makes it faster and easier to secure permission for the reuse of Springer Nature content to be published, for example, in a journal/magazine, book/textbook, coursepack, thesis/dissertation, annual report, newspaper, training materials, presentation/slide kit, promotional material, etc.

Simply visit SpringerLink and locate the desired content;

- Go to the article or chapter page you wish to reuse content from. (Note: permissions are granted on the article or chapter level, not on the book or journal level). Scroll to the botton of the page, or locate via the side bar, the "Reprints and Permissions" link at the end of the chapter or article.
- Select the way you would like to reuse the content;
- Complete the form with details on your intended reuse. Please be as complete and specific as possible ao as not to delay your permission request;
- Create an account if you haven’t already. A RightsLink account is different than a SpringerLink account, and is necessary to receive a licence regardless of the permission fee. You will receive your licence via the email attached to your RightsLink receipt;
- Accept the terms and conditions and you’re done!

For questions about using the RightsLink service, please contact Customer Support at Copyright Clearance Center via phone +1-855-239-3415 or +1-978-646-2777 or email springernaturesupport@copyright.com.

How to obtain permission to reuse Springer Nature content not available online on SpringerLink

Requests for permission to reuse content (e.g. figure or table, abstract, text excerpts) from Springer Nature publications currently not available online must be submitted in writing. Please be as detailed and specific as possible about what, where, how much, and why you wish to reuse the content.

Your contacts to obtain permission for the reuse of material from:

- books: bookpermissions@springernature.com
- journals: journalpermissions@springernature.com

Author reuse

Please check the Copyright Transfer Statement (CTS) or Licence to Publish (LTP) that you have signed with Springer Nature to find further information about the reuse of your content.

Authors have the right to reuse their article's Version of Record, in whole or in part, in their own thesis. Additionally, they may reproduce and make available their thesis, including Springer Nature content, as required by their awarding academic institution. Authors must properly cite the published article in their thesis according to current citation standards.
Material from: 'AUTHOR, TITLE, JOURNAL TITLE, published [YEAR], [publisher - as it appears on our copyright page]'

If you are any doubt about whether your intended re-use is covered, please contact journalpermissions@springernature.com for confirmation.

Self-Archiving

- Journal authors retain the right to self-archive the final accepted version of their manuscript. Please see our self-archiving policy for full details: https://www.springer.com/gp/open-access/authors-rights/self-archiving-policy/2124

- Book authors please refer to the information on this link: https://www.springer.com/gp/open-access/publication-policies/self-archiving-policy

The Geneva Papers on Risk and Insurance - Issues and Practice (2023) 48:502–547
https://doi.org/10.1057/s41288-023-00289-7

Risk mitigation services in cyber insurance: optimal contract design and price structure

Gabriela Zeller[1] · **Matthias Scherer**[1]

Received: 14 April 2022 / Accepted: 30 January 2023 / Published online: 8 May 2023

Abstract

As the cyber insurance market is expanding and cyber insurance policies continue to mature, the potential of including pre-incident and post-incident services into cyber policies is being recognised by insurers and insurance buyers. This work addresses the question of how such services should be priced from the insurer's viewpoint, i.e. under which conditions it is rational for a profit-maximising, risk-neutral or risk-averse insurer to share the costs of providing risk mitigation services. The interaction between insurance buyer and seller is modelled as a Stackelberg game, where both parties use distortion risk measures to model their individual risk aversion. After linking the notions of pre-incident and post-incident services to the concepts of self-protection and self-insurance, we show that when pricing a single contract, the insurer would always shift the full cost of self-protection services to the insured; however, this does not generally hold for the pricing of self-insurance services or when taking a portfolio viewpoint. We illustrate the latter statement using toy examples of risks with dependence mechanisms representative in the cyber context.

Keywords Cyber risk · Cyber insurance · Cyber assistance · Prevention · Self-protection · Self-insurance · Coherent risk measures · Stackelberg game

✉ Gabriela Zeller
gabi.zeller@tum.de

Matthias Scherer
scherer@tum.de

[1] Fakultät für Mathematik, Lehrstuhl für Finanzmathematik, Technische Universität München, Parkring 11, 85748 Garching, Germany

Introduction

Motivation and approach

Cyber insurance is still a relatively new, but steadily expanding market. The reasons for its ongoing growth in demand are manifold: the dynamically expanding and evolving cyber-threat landscape (ENISA 2021; tenable 2021), extensive media coverage of severe cyber incidents (Advisen and PartnerRe 2017, 2018; Marotta et al. 2017), ubiquitous introduction of stricter legislation (Anchen and Pain 2017; Marotta et al. 2017), and increased awareness of companies about their augmented dependence on information technology. To emphasise the first point, in particular the growing extent of the professionalism and economic potential of the ransomware "industry" are addressed, e.g. in ENISA (2021). As of 2020, cyber incidents were ranked the number one peril to businesses worldwide (Allianz 2020) and their perilousness can hardly be expected to have diminished since, as the COVID-19 pandemic and its effects (e.g. extensive ad-hoc shifts to remote work without adequate time to amend IT security measures and practices) have been labelled by some experts "the largest-ever cybersecurity threat" (Munich Re 2021). Many insurers are already actively participating in the global cyber insurance market, while still grappling with a firm understanding of this new and dynamic type of risk and its underlying drivers. Far from being solved is the question of how to adequately assess and price cyber risk given the various challenges, e.g. scarcity of historical data, non-stationarity of claims, association between claims, and strategic motivations of threat actors. Many academic works have recently been devoted to understanding and modelling these challenges in cyber risk. We, therefore, deliberately refrain from providing an exhaustive overview and refer to the surveys (Marotta et al. 2017; Awiszus et al. 2023).

In most established insurance lines, insurers have multiple years of claims experience and established technical expertise to quantify risks. In contrast, assessing and pricing cyber risks is particularly challenging due to the dynamically evolving threat landscape and the high complexity of modern IT systems. Therefore, insurers strive to collaborate with specialised IT security service providers (consider Bosch CyberCompare as an example or Advisen for a market overview), who not only support insurers in accurately assessing to-be-insured risks, but collaborate in providing services that aim at mitigating the insured risk as part of an insurance policy. Such cyber-assistance services can be divided into *pre-incident* services, such as network security, back-up of critical systems and data, and patch management, and *post-incident* services, such as restoration of data, forensic services, and legal advice (see Munich Re 2021). The former typically serve to decrease the probability of a cyber incident, while the latter support mitigation of the loss size in case an incident has occurred. In practice, the effects of both types of service are naturally intertwined, and additionally, all types of cyber assistance can also serve to provide insurers with additional information, i.e. to enhance their cyber-risk assessment practices or simply to obtain supplementary data (see also Remark 1 below). A recent survey (Munich Re

2021) indicates that the majority of (prospective) buyers believes that such services should be covered by holistic cyber insurance solutions, indicating that both the supply and demand side have realised that cyber insurance coverage should encompass more than pure compensation for financial losses. One type of service which is not yet explicitly advertised on the market, but holds great potential, is the insurer's ability to use the interdependence of cyber incidents to all parties' benefit by offering additional risk mitigation services.

To the best of our knowledge, established actuarial pricing approaches for these new policies are yet to be developed. The aim of this work is to propose a mathematical framework to study the optimal price structure of such insurance contracts, in particular to start addressing the question if (and under which circumstances) an insurer is economically incentivised to subsidise risk reduction services within an insurance policy. As part of this question, the issue of the optimal combination of insurance and risk mitigation (depending on their prices) from an insurance buyer's point of view is also studied. A further point, which is particularly relevant in the cyber context, is that for an insurer, it is not exhaustive to consider every single policyholder separately, but due to the potential interconnectedness of cyber losses, a portfolio viewpoint considering dependencies needs to be taken into account.

Our approach is based on the work of Bensalem et al. (2020), by using the framework of distortion risk measures and stochastic ordering of loss distributions, respectively, to capture risk assessment of all parties and the effects of risk mitigation services, and by modelling the interaction between insurer and insurance buyer(s) as a Stackelberg game. We extend their setting to a bivariate problem for the insurer, allowing her to choose the price for both risk transfer and risk mitigation, and analyse the results of the corresponding buyer's problem [which is conceptually similar to Bensalem et al. (2020)] in the cyber insurance context. Furthermore, we transcend from the study of an interaction with a single buyer to examples of (sequential or simultaneous) interactions with several buyers with dependent losses.

Related literature

A concise overview of academic studies on the interaction between risk reduction and insurance in the cyber context is given in Xiang et al. (2021). As mentioned therein, many of these studies rely on very simplified assumptions regarding the distribution of random cyber losses or the interplay between costs of prevention and consequence on the reduction of risk. Most often, the optimal combination of security provisions and insurance from an insured's point of view is studied, see, e.g. the early game-theoretic contribution of Pal and Golubchik (2010), the work of Young et al. (2016), and subsequently Mazzoccoli and Naldi (2020), or Yang and Lui (2014), Chase et al. (2017), and Mazzoccoli and Naldi (2021) who investigate optimal security investments under the presence of cyber insurance in a heterogeneous network, in a cloud computing environment, and for a multi-branch firm with correlated vulnerabilities, respectively. Zhang and Zhu (2021) use a dynamic moral hazard type of principal–agent model with Markov decision processes to capture decisions on self-protection of the insured and Skeoch (2022) expands the Gordon–Loeb model (Gordon and Loeb 2002) for cybersecurity to a cyber insurance context. Pal et al. (2014, 2017) more generally study synergies between cybersecurity and the (existence of a then nascent) cyber insurance market.

Fewer studies emphasise the insurer's role in designing cyber insurance contracts, e.g. by choosing premium and contractual indemnity (Dou et al. 2020), employing a bonus-malus system (Xiang et al. 2021), or trying to mitigate moral hazard by means of risk preference design (Liu and Zhu 2022).

The problem of combining different strategies of coping with risk, in particular the combination of risk reduction by investing in prevention measures and risk transfer by purchasing insurance, is of course not specific to cyber and has been the interest of many earlier studies. A good starting point is the survey (Courbage et al. 2013) on the economic literature on prevention and precaution. As differentiated therein, prevention activities encompass *self-protection*, i.e. modifying the probability of a loss, and *self-insurance*, i.e. shaping the potential loss size. The seminal work by Ehrlich and Becker (1972) examined the relationship of both activities to market insurance, and many authors have subjected these results to various model changes (for an overview, see Courbage et al. 2013), see, e.g. Dionne and Eeckhoudt (1985) and Hiebert (1989). Most aforementioned models use an *Expected Utility* (EU) framework and consider only two states (i.e. a loss occurs = "bad" state or no loss occurs = "good" state).[1] Another model of behaviour under risk, namely *Rank Dependent Expected Utility* (RDEU), has been considered for the study of prevention, e.g. in Konrad and Skaperdas (1993), Bleichrodt and Eeckhoudt (2006), Etner and Jeleva (2013). Courbage (2001) considered the relationships between market insurance, self-insurance, and self-protection in the context of Yaari's Dual Theory.

Our work is conceptually most closely related to Bensalem et al. (2020), who model the interaction between insurer and insurance buyer as a so-called *Stackelberg game* (see, e.g. Osborne and Rubinstein 1994; Fudenberg and Tirole 1991), a setting recently used to describe the interaction between reinsurer(s) and insurer(s), e.g. in Bai et al. (2022), Chen and Shen (2018), Chen et al. (2020), and Cheung et al. (2019).[2] Recently, some authors have also studied equilibria in sequential optimisation games in an insurance-reinsurance-setting, see, e.g. Boonen and Ghossoub (2022), Boonen et al. (2021) and Boonen and Zhang (2022). Let us also mention that in the cyber insurance domain, some works employ different game-theoretic approaches including the insurer and insured as parties, sometimes additionally featuring malicious third parties (cyber attackers), see, e.g. Zhang et al. (2017) and Yin et al. (2021). One aspect of the usual (principal-agent)problem between an insurer (acting as principal) and an insurance buyer (responding as agent) is the problem of *moral hazard*, i.e. the fact that the (risk reduction) actions of the agent are unobservable to the principal (see, e.g. Holmstrom 1979). This complicates matters, i.e. static principal-agent problems involving moral hazard are typically hard to solve (see,

[1] The distinction between self-protection and self-insurance provides good guidance and fits well with simple two-state models and frequency-severity approaches. Note that in reality, the effects of altering loss probabilities and loss sizes are often inseparable, which may be particularly relevant for cyber risks.

[2] The cited studies use a continuous-time setting, whereas we consider a one-period model.

e.g. Rogerson 1985; Jewitt 1988). Many of the above-mentioned works incorporate, or at least mention, the issue of asymmetric information in their studies, e.g. Liu and Zhu (2022), Boonen et al. (2021), and Zhang and Zhu (2021).[3]

The popular framework of risk measures to model risk preferences of both the insurance buyer and insurer has recently been used by, e.g. Bensalem et al. (2020), Cheung et al. (2019), Boonen and Ghossoub (2022), and Balbás et al. (2011), mostly in an insurance-reinsurance context. In the insurance context, an axiomatic characterisation of insurance prices as Choquet integrals (see Denneberg 2013) with respect to distorted probabilities was introduced in Wang et al. (1997) and studied further, e.g. in Bellini and Caperdoni (2007) and Wang (2000).[4] The first explicit connection of distortion risk measures and insurance pricing was made by introducing the *proportional hazard transform* (Wang 1995, 1996, 1998). Wang et al. (1997) described an axiomatic characterisation of insurance prices as Choquet integrals and Wang (2000) introduces another particular distortion in the general setting of Wang (1996), later called *Wang transform*, with the aim of connecting the pricing of insurance and financial risks. Finally, let us mention that many questions that arise from the practical usage (due to corresponding regulatory frameworks) of the value-at-risk (VaR) and average value-at-risk (AVaR) measures are subsequently studied for a more general class of distortion risk measures, e.g. backtesting methods [see, e.g. Christoffersen and Pelletier (2004) and Ziggel et al. (2014) for VaR, Emmer et al. (2015) and Kratz et al. (2018) for AVaR, and Bettels et al. (2022) for general distortion risk measures and an extensive overview of works on VaR and AVaR backtesting] or risk sharing [see, e.g. Galchion (2010) for VaR, Embrechts et al. (2018) for quantile-based risk measures (range value-at-risk), and Wang (2016), resp. Weber (2018), for more general (resp. *VaR-type*) distortion risk measures].

[3] Indeed, in other insurance domains, if incentive programmes exist (e.g. discounts on health insurance for participating in fitness regimes), they often give rise to moral-hazard issues, i.e. the insurer needs to secure the insured actually complies with the agreed-upon level of effort. In the cyber context, however, moral hazard does not seem to be a major concern for two reasons: first, due to the novelty and dynamics of cyber risk and the high complexity of technical systems, it is likely that neither of the parties (insurer and insured) have a full understanding of the underlying risk, i.e. the main problem is a lack of information for both parties rather than information being withheld. Due to the necessity for up-to-date technical expertise, insurers collaborate with specialised IT service providers to assess and monitor the insured risks and recommend or employ risk mitigation measures. Thus, in our framework, we assume both risk transfer and risk reduction are offered through the insurer (principal), i.e. risk reduction services are part of the insurance contract and therefore their uptake (ex-ante) and upholding (ex-post) observable to the insurer. Second, as e.g. reputational risk from cyber events or losses from threats classified as war actions are not fully insurable but substantial risks in practice, the insured has an intrinsic motivation to mitigate such risks, even if an insurance policy to transfer other financial losses is in place.

[4] Such *distortion risk measures* result from the properties of law-invariant, coherent risk measures if the property of sub-additivity for all random variables is replaced by additivity for comonotone random variables (see, e.g. Föllmer and Schied (2016) and Dhaene et al. (2012) for a detailed exposition and Dhaene et al. (2006, 2011) for a general review on (distortion) risk measures and their relation to comonotonicity). The sub-class of distortion risk measures with concave distortion functions used in this study can furthermore be shown to be coherent (see Wirch and Hardy 1999), i.e. are a sub-class of law-invariant, coherent risk measures.

Contribution

This paper extends the landscape of previous studies on the combination of risk reduction and risk transfer by bestowing the insurer with a more central role, namely controlling the cost of both risk transfer and risk mitigation. This relates to the real-world situation in cyber insurance, where insurers have started to endow insurance policies (risk transfer) with so-called cyber-assistance services (risk mitigation). We consider a monopolistic, profit-maximising, risk-averse or risk-neutral insurer using a concave distortion risk measure and study separately the cases of cyber-assistance services relating to the concepts of self-protection and self-insurance.[5] The interaction between the insurer and the insurance buyer(s),[6] who are risk averse and also use a concave distortion risk measure, is modelled as a Stackelberg game, where the "inner" optimisation problem corresponds to the insurance buyer's response to a given price structure by the insurer and the "outer" optimisation problem corresponds to the insurer's problem of determining prices for (cyber) risk transfer and (cyber) assistance services. In particular, we derive the following insights:

- The "The insurer's problem: single-contract case" section addresses the insurer's problem in the single-contract case, studying in which cases an insurer is incentivised to encourage risk reduction in her policyholders by sharing the cost of risk reduction measures. We find that under the above assumptions, the insurer would never share the cost of risk reduction in a single-contract, pure self-protection scenario (Theorem 1 and case study in section A.5 in the electronic supplementary information). This does not hold in a single-contract, pure self-insurance scenario, where the optimal share of risk mitigation cost the insurer chooses to bear may depend e.g. on the parameters of the loss size distribution and both parties' risk aversions (Remark 11 and case study in section A.6 of the electronic supplementary information).
- The "The insurer's problem: portfolio viewpoint" section extends the insurer's study of the pure self-protection scenario from a single-contract view to bivariate examples of insurance buyers facing dependent cyber losses under dependence mechanisms relevant for cyber (loss propagation, common events). We demonstrate that the finding from the single-contract case does not carry over, i.e. already for these small toy portfolios, the insurer may have an incentive to subsidise risk mitigation in some policyholders. The study is extended to an example of a larger ($N \geq 2$) portfolio in section A.7.3 in the electronic supplementary information, illustrating the increasing importance of taking a portfolio viewpoint for dependent risks.

[5] While both types of services can have intertwined effects and relate to gaining information via risk assessment services, the issues of moral hazard / asymmetric information and the prospect of gaining additional information are excluded from the mathematical analysis in the main part of this paper. A discussion of how to potentially address the effect of risk assessment services is provided in section A.1 in the electronic supplementary information.

[6] We consider a single buyer during the first part of the paper and extend this to examples of two (resp. $N \geq 2$) buyers with dependent cyber-loss occurrences in the "The insurer's problem: portfolio viewpoint" section (resp. section A.7.3. in the electronic supplementary information).

- The "Solution to the insurance buyer's problem" section addresses the insurance buyer's solution to his problem of choosing an optimal combination of insurance and risk mitigation for a given price structure by the insurer (Corollary 3) and deduces the potentially complementary nature of the two activities (Corollary 4).

In summary, the contribution offers threefold insights, regarding the viewpoints of insurers, (prospective) insurance buyers, and the general (cyber insurance) market. For insurers, the study of the insurer's bivariate optimisation problem offers a first guidance to the optimal pricing of insurance policies including risk mitigation services (under specific assumptions). For insurance buyers, it is also invaluable to better understand how different contracts would be optimally priced by an insurer, in particular that the price structure a prospective policyholder is offered (and the included incentive for risk reduction) may not only depend on his own characteristics, but on the insurer's existing portfolio and the (assumed prospective) dependence between losses.[7] The study of the insurance buyer's problem on the optimal combination of risk transfer and risk mitigation is not conceptually new, but its detailed consideration offers valuable insights. Next to naturally providing guidance on the recommended course of action for insurance buyers, it may serve to theoretically explain the *insurance gap* observed in the cyber insurance market (see, e.g. Shetty et al. 2018), an offer-demand mismatch caused by the fact that potential buyers often look for insurance against extreme cyber events and tend to perceive asked prices of such coverage as excessive, while insurers seek to limit their liabilities from unprecedented cyber losses either by limiting coverage or by charging heavy risk premiums. One way to mitigate this mismatch, where no premium acceptable to both parties can be found for the original risk, is to equip insurance policies with (potentially subsidised) risk reduction services which help to alter the risk in a way that allows the insurer to reduce premiums and offer desired coverage at an acceptable (from the buyer's viewpoint) premium.

The remainder of this paper is structured as follows: in the "Model set-up and assumptions" section, the model assumptions and set-up are explained; in the "Solution to the insurance buyer's problem" and "The insurer's problem: single-contract case" sections the insurance buyer's and insurer's optimisation problems, respectively, are studied in the single-contract setting; the "The insurer's problem: portfolio viewpoint" section addresses the insurer's problem in simple portfolio settings with dependent losses. The "Conclusion" section summarises and outlines future research opportunities.

[7] This implies that a prospective buyer would be particularly well advised to enquire about prices at different insurers, as the offered price structures may differ depending on the existing portfolio, even if (hypothetically) the insurers' risk assessment and modelling processes were identical.

Model set-up and assumptions

Risk mitigation services in cyber insurance (cyber assistance)

We first consider a model involving one profit-maximising, risk-averse insurer ('she') and one risk-averse (insurance) buyer ('he'). Before detailing the model set-up and the mechanics of the sequential optimisation game, we give some compelling arguments for considering *risk mitigation services* in conjunction with cyber insurance policies and subsume types of risk mitigation services into three categories:

(R1) *Reduction of loss probability after initial risk assessment:* Insurers often work with specialised IT service providers (SP) who help them to thoroughly classify a prospective client's IT security. After the effort of such an assessment is invested, the SP and the assessed company share a common understanding of the company's IT security standpoint and potential need for action. Given that the risk is deemed insurable, a joint offer by SP and insurer to the company is in everyone's interest: the company receives insurance protection and high-quality IT security maintenance services as a joint package without the necessity of extra effort to ensure complying with the insurer's requirements, which is especially relevant for small companies. The insurer does not forfeit the upfront investment for risk assessment and has certainty about the maintenance and potential improvement of the IT security according to the SP's assessment. The SP has certainty about the company's willingness to comply with recommendations in order not to jeopardise insurance coverage, and about insurance coverage with a trusted "counterparty" who will not doubt their work in case a cyber event still occurs.[8]

(R2) *Reduction of loss magnitude in a cyber event:* Among the insured's obligations within a typical cyber insurance contract is the immediate notification of the insurer in case of a (suspected) cyber event. This allows the insurer to supply immediate technical and legal support in order to mitigate economic losses. Naturally, it is in both the company's and insurer's interest for these experts to already have a good understanding of the company's IT security landscape and to be available immediately, both of which can be guaranteed by including these services – to be performed by a service provider collaborating with the insurer – in an insurance contract.

(R3) *Use of insurer's knowledge about current cyber-loss landscape:* While many businesses dedicate their attention to describing current cyber-threat trends, insurers have invaluable knowledge about economic losses currently suffered by their portfolio of clients. Companies are usually obliged by contract to notify their cyber insurer about cyber events, while naturally being reluctant to voluntarily share this information publicly or with external parties (e.g. researchers) in order to avoid reputational damage. Therefore, insurers have an information

[8] All of the above considerations emphasise again that moral hazard and information asymmetries might not be a severe problem in cyber as knowledge and incentives are aligned.

advantage regarding current threats and their common causes (e.g. a new trend in phishing mails or a vulnerability in a software used by companies of a specific industry sector) and can make use of this extra knowledge to warn other policyholders who are particularly prone to similar threats and vulnerabilities (e.g. all policyholders from the same industry sector or all using some vulnerable software). The benefit of doing so is reducing the probability of additional cyber losses from the same cause in their portfolio. This is especially relevant for large companies with sophisticated IT security (who may already work with external SPs) which might not find it necessary to additionally take advantage of (R1) and (R2) as part of insurance coverage. For the insurer, this type of mitigation helps to reduce the impact of systemic events and, thus, accumulation risk in the portfolio.

Remark 1 (Link between theoretical and marketed types of risk reduction service) The types of service currently offered on the cyber insurance market and suggested above direct quite naturally to the concepts of self-protection and self-insurance:

(R1) Describes *pre-incident services* which are *self-protection activities*. Examples are network security, back-up of critical systems and data, anti-malware tools, identity and access management, IT security consulting, employee awareness measures, patch management, and mobile device management (Munich Re 2021).

(R2) Describes *post-incident services* which are *self-insurance activities*, such as restoration of data, 24h help hotlines, forensic post-breach services, legal advice, and consulting in case of extortion (Munich Re 2021).

(R3) Describes a type of *self-protection activity* not yet advertised on the market, as contracts are typically viewed stand alone. However, using the insurer's portfolio knowledge to install such warning mechanisms would be an important way to use dependencies (and information) between risks to the insurer's and insureds' advantage.

Of course, the above categorisation simplifies reality regarding several points: pre- and post-incident services are usually not offered disjointly, but as a complete "cyber assistance" service package, and each service activity within the above categories can have beneficial effects on both cyber-loss probability and severity. For example, anti-malware tools not only serve their primary purpose, i.e. to deter malware from entering the system (preventing a cyber incident completely), but as a side effect – in case malware circumvents the protection – may help to identify the source of a cyber incident more efficiently and reduce the time until system functionality is restored (reducing the economic impact of an occurred cyber incident). Nevertheless, from a mathematical viewpoint, it is convenient (and in line with previous academic work) to study the two concepts separately and therefore it is helpful to keep in mind the types of "real-world cyber assistance activities" they relate to.[9] One aspect of cyber assistance which is purposely omitted here is risk-assessment services (see section A.1 in the electronic supplementary information). This includes, e.g. extensive IT audits conducted by an IT service provider collaborating with the insurer to analyse a company's IT security provisions, to identify vulnerabilities, and to provide recommended courses of action.

Model prerequisites

Following the framework of Bensalem et al. (2020), we assume that over a given policy year, the buyer faces a random loss represented by a non-negative random variable (r.v.) X from a family of distributions F_s indexed by a parameter $s \in [0, \infty)$.[10] For $X \sim F_s$, we denote the corresponding survival function by $\overline{F}_{X,s}(x) = \mathbb{P}_s(X > x),\ x \in \mathbb{R}$, and its generalised inverse, the tail quantile function, by $\overline{q}_{X,s}(u) = \overline{F}^{-1}_{X,s}(u) = \inf\{x \in \mathbb{R} : \overline{F}_{X,s}(x) \leq u\},\ u \in (0, 1)$. To formalise the relationship between the parameter s and the distributions F_s, we assume a decreasing order in the sense of first-order stochastic dominance ($\leq_{FSD}$), i.e. for any $0 \leq s_1 < s_2 < \infty$ and $X_1 \sim F_{s_1},\ X_2 \sim F_{s_2}$ it holds that $X_2 \leq_{FSD} X_1$. This is equivalent (see Müller and Stoyan (2002), Theorem 1.2.8) to assuming

$$0 \leq s_1 < s_2 < \infty \implies \mathbb{E}_{s_2}\left[f(X)\right] \leq \mathbb{E}_{s_1}\left[f(X)\right]$$

for any non-decreasing[11] function $f : \mathbb{R} \to \mathbb{R}$ for which both expectations exist. We furthermore assume that $\mathbb{E}_s[X] > 0,\ \forall s \in [0, \infty)$, meaning that no risk reduction can ever completely eliminate the possibility of a positive loss.

The decreasing order in the sense of FSD of F_s implies that

$$\text{for any} \quad u \in (0, 1), \quad \text{the map} \quad s \mapsto \overline{q}_{X,s}(u) \quad \text{is non-increasing.} \tag{A1}$$

This means that increasing s alters the risk X in such a way that for any probability level, the minimum loss amount that is exceeded by X with this probability does not increase.

Assumption 1 (Convexity of tail quantile in s). Furthermore, we assume that

$$\text{for any} \quad u \in (0, 1), \quad \text{the map} \quad s \mapsto \overline{q}_{X,s}(u) \quad \text{is convex.} \tag{A2}$$

[9] Naturally, an extension to a setting where both concepts are studied as intertwined remains an interesting task for future research.

[10] The parameter s denotes the amount of risk mitigation service, whose categories were detailed above.

[11] Throughout, we use the term *non-decreasing* for a real-valued function that fulfils $\forall x, y : x < y \implies f(x) \leq f(y)$ and *increasing* if the order in the implication is strict. The terms *non-increasing* and *decreasing* are used analogously.

This assumption can be interpreted as a decrease in marginal effect of service, i.e. the impact per unit of s on the risk X in the sense of (A1) does not increase as the baseline level of s increases, which is a very natural economic assumption.

We assume that both parties evaluate risk by using law-invariant, coherent risk measures, whose properties are recalled in section A.2 of the electronic supplementary information. An important class of risk measures are so-called distortion risk measures (see Wang et al. 1997), defined for a real-valued r.v. X as the usual Choquet integral that simplifies for non-neg. X to

$$\rho(X) := \int_0^\infty \psi(\overline{F}_X(x))\mathrm{d}x \overset{e.g.\,[32]}{=} \int_0^1 \overline{q}_X(u)\mathrm{d}\psi(u), \tag{1}$$

where $\psi : [0,1] \to [0,1]$ is a distortion function[12] and $\overline{q}_X(u),\ u \in (0,1)$, is the tail quantile function. From Eq. (1), one can directly see that the distortion risk measure for a.s. non-neg. losses represents a distorted expectation of X.

Assumption 2 (Concavity of distortion function). Concavity of the distortion function is a natural economic assumption. As it corresponds to assigning a higher weight to small probability events, it describes risk aversion of the decision maker, a standard assumption and indeed a prerequisite for the existence of insurance. Therefore, we will restrict our analysis to distortion risk measures with concave distortion, a class of coherent, law-invariant risk measures.[13]

Remark 2 [Distortion risk measures and stochastic dominance, e.g. Dhaene et al. (2006)] Any distortion risk measure ρ preserves first-order stochastic dominance, i.e. for any a.s. non-negative r.v. X_1, X_2, it holds that $X_1 \leq_{FSD} X_2 \implies \rho(X_1) \leq \rho(X_2)$.

Example 1 Table 1 lists some commonly used distortion risk measures and their corresponding distortion functions. In the case studies of our latter analysis, we focus on the *proportional hazard transform*.

The above assumptions on the risk measures and loss distributions [in particular (A2)] are convenient insofar as they imply that the map $s \mapsto \rho_s(X)$ (and as a special case $s \mapsto \mathbb{E}_s[X]$) is convex, continuous, non-increasing, and $\rho_s(X) \geq \mathbb{E}_s[X] > 0$ [see Bensalem et al. (2020) and section A.3 in the electronic supplementary information].

[12] A distortion function $\psi : [0,1] \to [0,1]$ is a continuous, non-decreasing function with $\psi(0) = 0$ and $\psi(1) = 1$. The distortion is often economically interpreted as a subjective weighting of objective probabilities representing the decision maker's views or risk preference.

[13] By the properties of the Choquet integral (see Denneberg 2013), any distortion risk measure fulfils 1., 2., 4., and 5. in Definition 1 (Section A.2 of the electronic supplementary information) and additionally 3. if the distortion function ψ is concave (and the underlying probability space has no atoms), see, e.g. Wirch and Hardy (1999).

Table 1 Popular distortion risk measures (DRM) and underlying distortion functions

Risk measure	Distortion $\psi(u)$, $u \in (0,1)$	ψ concave	Parameters and remarks
VaR_α	$\mathbf{1}_{\{u>1-\alpha\}}$	No	$\alpha \in (0,1)$
$AVaR_\alpha$	$\min\left\{\frac{u}{1-\alpha};1\right\}$	Yes	$\alpha \in (0,1)$
Wang transform RM (Wang 2000)	$\Phi\left(\Phi^{-1}(u)+\lambda\right)$	Yes	$\lambda \in (0,\infty)$, Φ is std. Normal c.d.f.
Beta DRM (Wirch and Hardy 2000)	$\frac{1}{\beta(a,b)}\int_0^u t^{a-1}(1-t)^{b-1}\mathrm{d}t$	Yes	$0 < a \leq 1,\ b \geq 1$, $\beta(a,b) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$
Proportional Hazard (PH) transform RM (Wang 1995)	u^r	Yes	$r \in (0,1]$, Special case of Beta DRM

Interaction between cyber-insurance buyer and insurer

We now describe how the interaction between insurance buyer and insurer in the case of a cyber insurance contract is modelled as a *Stackelberg game*, i.e. a sequential optimisation game between two parties, where one party (the *leader*) moves first by choosing her strategy and the other party (the *follower*) moves second by choosing his strategy depending on the selected strategy of the leader, whereby both parties seek to maximise a gain or utility function or equivalently, minimise a loss function. For a general introduction to Stackelberg games, see Fudenberg and Tirole (1991) and Osborne and Rubinstein (1994). A common tool to solve a Stackelberg game is *backward induction* (see Fudenberg and Tirole 1991), i.e. first solving the follower's problem for any possible choice of the leader's strategy and then – knowing all the follower's responses – solving the leader's problem. The search for a solution (and its existence) therefore depends on the specific formulation of both problems, which we now detail in our case.

0. Common (correct) knowledge of initial loss distribution

The prospective insurance buyer approaches the insurer to inquire about offered prices for cyber insurance policies (in person or by entering data into an online calculation system), where in order to receive price quotes, he needs to provide information that allows the insurer (with the help of an IT service provider) to classify his risk profile given his characteristics (e.g. industry sector, company size, IT security measures). We assume he provides the information truthfully and to the best of his knowledge, such that buyer and insurer have a common, unambiguous view of the original loss distribution, denoted F_0.[14] The real-world uncertainty of either

[14] F_0 denotes the loss distribution of the buyer given his initial characteristics, including his existing IT security measures. The subscript 0 indicates that no *additional* services to reduce the risk have yet been acquired following the initial risk assessment. As the initial IT security level (and other characteristics) vary between prospective buyers, the initial risk assessment yields inhomogeneous F_0. Note that for some companies, the risk assessment as part of the insurance take-up process may be the first comprehensive analysis of the cybersecurity level of their organization. While not every inquiry about insurance prices leads to the closure of a cyber insurance contract, the process may serve as a wake-up call for the acquisition of (additional) risk reduction measures within or without an insurance policy.

parties' knowledge of the unknown initial loss distribution is not studied here. Naturally, the question of accurate cyber-risk assessment has gained increased practical importance and expresses itself, e.g. in the increasing number of service providers in this domain, see, e.g. Bosch CyberCompare as an example or Advisen for a market overview. For a seminal discussion of cyber-risk assessment services and a proposal how to approach them mathematically, see section A.1 in the electronic supplementary information.

1. Prices quotes by the insurer

Given the buyer's original risk $X \sim F_0$, the insurer offers price quotes Π for a range of contracts, where each offered contract is characterised by the included level of *risk mitigation service* $s \in [0, \infty)$.[15] Assume that the price of entering a contract with service level $s \in [0, \infty)$ is given by

$$\Pi(s) = (1 + \theta)\mathbb{E}_s[X] + \beta c(s),$$

where the first term represents the *risk premium* according to the expected value principle with loading θ and the second term denotes the *service premium*, where we assume that providing service at level $s \in [0, \infty)$ requires a monetary cost of $c(s)$ for the insurer, of which a proportion $\beta \in [\underline{\beta}, 1]$, $\underline{\beta} > 0$,[16] is charged to the insured and, thus, the remaining proportion $(1 - \beta)$ can be regarded a subsidy by the insurer to incentivise risk reduction. Analogously to (A1) and (A2), $s \mapsto c(s)$ is assumed to be increasing, strictly convex, and continuous with $c(0) = 0$ and $\lim\limits_{s \to \infty} c(s) = \infty$. The cost incurred by the insurer can be understood e.g. as the internal cost charged by the IT service provider for providing pre- or post-incident services (i.e. (R1) and (R2)) or the administrative cost of monitoring and evaluating loss data to warn policyholders about imminent threats (i.e. (R3)). Thus, the insurer's task is to choose a combination $(\theta, \beta) \in [0, \infty) \times [\underline{\beta}, 1]$ which then defines price quotes for all feasible contracts.

2. Choice of a contract by the buyer (or opt-out)

Given a family of prices $\Pi(s)$ for all feasible contracts, the buyer selects a contract by choosing a proportional insurance share $\alpha \in \{0, 1\}$ (to opt into full insurance $\alpha = 1$ or to not buy insurance $\alpha = 0$) and the amount of risk mitigation service $s \in [0, \infty)$. We assume that the purchase of (additional) service at any level s is also feasible outside of an insurance contract, but at a higher cost $\beta_o c(s)$ with $\beta_o > 1$. This can be understood as the cost of buying service directly through an IT service provider (without a discount offered for insurance customers) or from the

[15] One might argue that s should rather be chosen from a discrete set $\{s_1, \dots, s_n\}$, $n \in \mathbb{N}$ (a potentially interesting combinatorial optimisation problem), representing all feasible combinations of service packages offered by the insurer. This is reasonable and we regard this as a mathematically different version of the problem whose analysis is not the present focus.

[16] As β does not depend on s (a potential generalisation for future studies), we do not allow the insurer to give away service for free, as otherwise the cost of service $\beta c(s)$ would not increase with its amount, which is unnatural.

insurer herself at a mark-up.[17] In summary, given the prices for all feasible contracts as offered by the insurer, the insurance buyer's problem consists of choosing $(\alpha, s) \in \{0, 1\} \times [0, \infty)$. We detail in Remark 4 how the insurance buyer's choice encapsulates three classical ways of dealing with risk (acceptance, reduction, transfer), see, e.g. Marotta et al. (2017).

3. Solution by backward induction

To find both parties' optimal solution, we use backward induction (see, e.g. Osborne and Rubinstein 1994) by first finding the buyer's optimal response (α^*, s^*) to any insurer's choice of (θ, β) and second, given all optimal buyer's responses, finding the insurer's optimal choice $(\theta^*(\alpha^*, s^*), \beta^*(\alpha^*, s^*))$. In order to formulate and solve the game, below we state the loss functions of buyer and insurer, respectively.

Remark 3 We highlight some similarities and distinctions between the present work and the study of Bensalem et al. (2020), whose framework was our inspiration: as indicated above, the choice of risk measures and the ordering of loss distributions follows Bensalem et al. (2020) and from the insurance buyer's point of view, the risk reduction service s fulfils a very similar role to the *effort* considered in Bensalem et al. (2020), yielding related optimisation problems for the buyer within the Stackelberg game. In the present study, however, the insurer's role is more central, as she controls the cost of risk mitigation service within an insurance contract (via the share β of administrative cost charged to the insured). This implies that the insurer has to solve a two-dimensional problem (choosing a combination of *risk premium* and *service premium* optimally), and circumvents the moral-hazard problem that often occurs in studies on prevention and insurance. As in the present setting the risk mitigation service is offered through the insurer, the challenge of ensuring that the buyer actually complies with the agreed-upon optimal level of risk reduction (according to which insurance is priced) does not arise. Furthermore, we extend the study of the interaction with one insurance buyer to toy examples of interactions with a portfolio of dependent buyers, a particularly relevant issue in the cyber context.

Formalisation of the Stackelberg game

We now combine the assumptions of the above sections to formulate the optimisation problems of both parties within a Stackelberg game. For the reader's convenience, all parameters and functions appearing within the optimisation problems are summarized in Tables 2 and 3. The insurance buyer's objective is to minimise a coherent and law-invariant risk measure ρ_1 associated to his total position including insurance, while the insurer's objective is to minimise, given the buyer's optimal response, another coherent and law-invariant risk measure ρ_0 associated to her (negative) total loss.

[17] The latter option is not necessarily feasible in practice, as the insurer may not be interested in or legally allowed to sell such services.

Table 2 Overview of parameters in problems (BP) and (IP)

Parameter	Range	Interpretation	Buyer's view	Insurer's view	Comment
α	$\{0, 1\}$	Insurance share (opt-in (= 1) / opt-out (= 0))	Optimisation parameter, choose α^*	α^* chosen by buyer	
s	$[0, \infty)$	Cyber service level	Optimisation parameter, choose s^*	s^* chosen by buyer	
$s_B(\theta, \beta)$	$(0, \infty)$	Boundary between sets of values of s where no insurance vs. full insurance is preferred in the self-protection case	Fixed	Fixed given choice of (θ, β)	See Corollary 1 and Remark 6.
$s_{B1/B2}(\theta, \beta)$	$(0, \infty)$	Boundaries between sets of values of s where no insurance vs. full insurance is preferred in a special self-insurance case	Fixed	Fixed given choice of (θ, β)	See section A.6 in the electronic supplementary information.
s_N	$[0, \infty)$	Optimal service demand without insurance	To be found by minimising $L_{1,\mathcal{N}}(s)$ (loss without insurance)	Fixed	See Corollary 2 and Remark 7.
$s_I(\theta, \beta)$	$[0, \infty)$	Optimal service demand with insurance	To be found by minimising $L_{1,\mathcal{I}}^{\theta,\beta}(s)$ (loss with insurance)	Fixed given choice of (θ, β)	See Corollary 2 and Remark 7.
θ	$[0, \infty)$	Risk loading	θ^* chosen by insurer	Optimisation parameter, choose θ^*	
θ_0	$[0, \infty)$	Minimum loading s.t. buyer would choose not to insure original risk	Fixed, constant	Fixed, constant	See Corollary 1 and Remark 6.
$\theta_N(\beta)$	(θ_0, ∞)	Minimum loading s.t. global minimiser of no-insurance problem s_N lies in set where no insurance is chosen	Fixed	Fixed given choice of β	See Corollary 2 and Remark 7.

Table 2 (continued)

Parameter	Range	Interpretation	Buyer's view	Insurer's view	Comment
$\theta_I(\beta)$	$[0, \infty)$	Maximum loading s.t. pure risk transfer (insurance) is preferred to combination of risk transfer and risk reduction (insurance and service)	Fixed	Fixed given choice of β	See Corollary 2 and Remark 7.
$\theta_R(\beta)$	$(\theta_N(\beta), \infty)$	Maximum loading s.t. buyer chooses (full) insurance	Fixed	Fixed given choice of β	See Corollary 3 and Remark 8; Denote $\underline{\theta} := \theta_R(1)$, $\bar{\theta} := \theta_R(\underline{\beta})$, see the "The insurer's problem: single-contract case" section.
β	$[\underline{\beta}, 1]$	Share of service cost shifted to buyer	β^* chosen by insurer	Optimisation parameter, choose β^*	
$\underline{\beta}$	$(0, 1)$	Minimum share of service cost to be shifted to buyer	Fixed, constant	Fixed, constant	Service cannot be given away for free.
β_o	$(1, \infty)$	Share of service cost without insurance	Fixed, constant	Fixed, constant	Service without insurance is more expensive than service combined with insurance.

Table 3 Overview of functions in problems (BP) and (IP)

Function	Interpretation	Properties	Comment
$\rho_{0,s}(X)$	Insurer's risk measure	$X \mapsto \rho_{0,s}(X)$ is a (coherent) distortion risk measure; $s \mapsto \rho_{0,s}(X)$ is convex, continuous, non-increasing, and $\rho_s(X) \geq \mathbb{E}_s[X] > 0$	See Definition 1 and Lemma 1.
$\rho_{1,s}(X)$	Insurance buyer's risk measure	As for $\rho_{0,s}(X)$	Risk measures should reflect that buyer is more risk-averse than insurer.
$c(s)$	Cost of service	$s \mapsto c(s)$ is increasing, strictly convex, and continuous with $c(0) = 0$ and $\lim_{s\to\infty} c(s) = \infty$	Increasingness and convexity are natural economic assumptions.
$p(s)$	Loss probability in the self-protection case	$s \mapsto p(s) \in [0, 1]$ is decreasing and convex	Decreasingness and convexity are natural economic assumptions. Additionally assume convexity of $s \mapsto \psi(p(s))$ to ensure convexity of problem (BP).

$$\min_{(\alpha,s)\in\{0,1\}\times[0,\infty)} L_1(\alpha,s) := \rho_{1,s}(X) + \beta_o c(s) + \alpha\Big[(1+\theta)\mathbb{E}_s[X] - \rho_{1,s}(X) + (\beta-\beta_o)c(s)\Big], \tag{BP}$$

$$\min_{(\theta,\beta)\in[0,\infty)\times[\underline{\beta},1]} L_0(\theta,\beta) := \alpha^*(\theta,\beta)\Big(\rho_{0,s^*(\theta,\beta)}(X) - (1+\theta)\mathbb{E}_{s^*(\theta,\beta)}[X] + (1-\beta)c(s^*(\theta,\beta))\Big), \tag{IP}$$

where we have used that both risk measures are cash-additive and positively homogeneous. It is obvious that the insurer's loss depends on (θ,β) directly as well as via the buyer's optimal response denoted $(\alpha^*(\theta,\beta), s^*(\theta,\beta))$.

Remark 4 (Interpretation of insurance buyer's choice) The buyer's options correspond to three classical ways of dealing with risk:

- **Risk acceptance:** The choice $(\alpha,s)=(0,0)$ yields $L_1(0,0)=\rho_{1,0}(X)$, i.e. is equivalent to opting out of buying insurance or services and just retaining and accepting the original risk.
- **Pure risk transfer:** Choosing $(\alpha,s)=(1,0)$ yields $L_1(1,0)=(1+\theta)\mathbb{E}_0[X]$, meaning that the buyer opts for fully insuring the original risk.
- **Pure risk reduction:** A choice $\alpha=0,\ s>0$ yields $L_1(0,s)=\rho_{1,s}(X)+\beta_o c(s)$, i.e. the buyer opts out of risk transfer but chooses to reduce the original retained risk by purchasing risk reduction services (from the insurer outside of a policy or from a service provider directly).[18]
- **Combination of risk transfer and risk reduction:** A choice $\alpha=1,\ s>0$ yields $L_1(1,s)=(1+\theta)\mathbb{E}_s[X]+\beta c(s)$ and means that the buyer chooses an insurance policy with risk mitigation services included, i.e. opts for insuring a reduced risk.

Remark 5 (Buyer's and insurer's optimal attainable loss)

- Note that as the insurance buyer starts out by facing the non-negative random loss X, by assumption $L_1(\alpha^*,s^*)>0$, i.e. the insurance buyer can never completely eliminate his risk or even make a profit.
- On the contrary, we naturally assume that the insurer only offers a contract if it is profitable, i.e. only if she can obtain a negative loss $L_0(\theta^*,\beta^*)<0$. Otherwise, she would refrain from offering a contract by refusing to quote a price.

[18] If one does not want to allow the interpretation that such contracts are offered by the insurer outside of an insurance policy (e.g. due to legal restrictions), the insurer's loss function should be formulated in a way that makes these contracts unprofitable (e.g. as done here by restricting $\beta\in[\underline{\beta},1]$). If one wants to allow such contracts (one could argue that such a contract could be closed in the cyber domain with a client that has other contracts with the same insurer), a choice of $\beta>1$ would allow the insurer to sell her services at a mark-up (one could argue that this might be profitable for an insurer who has the appropriate infrastructure in place anyway for the rest of her portfolio). In our analysis, we stick to the interpretation that these *outside* service contracts are offered by third parties, i.e. service providers, and their price is externally given and higher than any within-insurance price (i.e. $\beta_o>1$, see above).

Solution to the insurance buyer's problem

As the analysis of (BP) is an extension of Bensalem et al. (2020), this section focuses on the additions to their analysis originating from the new formulation of (IP) and the interpretation of all results in the cyber insurance context. Derivations and proofs are outlined in section A.3 of the electronic supplementary information. First, one determines the set of values of s such that full insurance is demanded (i.e. $\alpha^*(s) = 1$, denoted $\mathcal{I}$) and its complement (no insurance is demanded, $\alpha^*(s) = 0$, $\mathcal{N} := \mathcal{I}^{\mathrm{c}}$). Note that for fixed s, the choice $\alpha^* \in \{0, 1\}$ depends only on the sign of the expression in the last bracket of (BP) such that it follows:

$$\alpha^* = 1 \iff G^\beta(s) := \frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]} + (\beta_o - \beta)\frac{c(s)}{\mathbb{E}_s[X]} \geq (1+\theta) \quad \implies \mathcal{I} := \{s \in [0,\infty) : G^\beta(s) \geq (1+\theta)\}, \tag{2}$$

$$\alpha^* = 0 \iff G^\beta(s) < (1+\theta) \quad \implies \mathcal{N} := \{s \in [0,\infty) : G^\beta(s) < (1+\theta)\}. \tag{3}$$

On the sets $\mathcal{I}$ and $\mathcal{N}$, the buyer's loss function is a sum of convex functions:

$$L_{1,\mathcal{N}}(s) := \rho_{1,s}(X) + \beta_o c(s), \; s \in \mathcal{N},$$
$$L_{1,\mathcal{I}}^{\theta,\beta}(s) := (1+\theta)\mathbb{E}_s[X] + \beta c(s), \; s \in \mathcal{I}.$$

Therefore, one considers (BP) separately on $\mathcal{I}$ and $\mathcal{N}$ and compares the resulting local minima to obtain a global minimum. To this end, one first needs to study $\mathcal{I}$ and $\mathcal{N}$ for given (θ, β), i.e. the behaviour of $s \mapsto G^\beta(s)$ with respect to the threshold $(1+\theta)$. We know that by assumption and Lemma 1 (see section A.3 in the electronic supplementary information), $s \mapsto G^\beta(s)$ is continuous and its second summand $s \mapsto (\beta_o - \beta)\frac{c(s)}{\mathbb{E}_s[X]}$ is non-negative and increasing.[19] In this study, we consider two cases:

- **Self-protection:** In a self-protection scenario (Ehrlich and Becker 1972), i.e. if service only affects the probability of a loss, the map $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$ is monotone non-decreasing (see Bensalem et al. 2020, Lemma 3.2, and section A.3 in the electronic supplementary information). Economically, this means that increased risk reduction has a larger impact on (reducing) the price of insurance than on (reducing) the risk.[20] Mathematically, this implies increasingness of the entire map

[19] This follows immediately as by assumption $\mathbb{E}_s[X] > 0$, $\beta_o - \beta > 0$, and $s \mapsto c(s)$ is non-negative and increasing, while by Lemma 1, $s \mapsto \mathbb{E}_s[X]$ is non-negative and non-increasing.

[20] This can be seen even more clearly by rewriting Equation (A3) in terms of *elasticity* η with respect to s [as used in economics for e.g. the price-elasticity of demand, see e.g. Parkin et al. (2002)], i.e. for $0 < s_1 < s_2 < \infty$ as

$$\underbrace{\frac{\frac{\mathbb{E}_{s_2}[X]-\mathbb{E}_{s_1}[X]}{\mathbb{E}_{s_1}[X]}}{\frac{s_2-s_1}{s_1}}}_{\eta_{\mathbb{E},s}} \leq \underbrace{\frac{\frac{\rho_{s_2}[X]-\rho_{s_1}[X]}{\rho_{s_1}[X]}}{\frac{s_2-s_1}{s_1}}}_{\eta_{\rho,s}} \leq 0,$$

yielding that the expectation is more elastic with respect to service than the risk measure.

$s \mapsto G^{\beta}(s)$, meaning that $G^{\beta}(s)$ could intersect (for given β and θ) the threshold $(1+\theta)$ at most once, making $\mathcal{I}$ and $\mathcal{N}$ straightforward to determine. This setting will be considered in the following.

- **Special case of self-insurance:** Bensalem et al. (2020) argue that in a scenario of self-insurance, i.e. in the present context if service only affects the severity of a cyber loss, for some standard loss distributions (e.g. Pareto, Weibull, or Log-Normal), $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$ is monotone non-increasing. This does not lead to a straightforward expression of $\mathcal{I}$ and $\mathcal{N}$, as monotonicity of $s \mapsto G^{\beta}(s)$ is not implied and there is a priori no limit for the number of times it crosses a given threshold $(1+\theta)$ for $s \in [0,\infty)$, such that no general results for this case can be stated. In section A.6 in the electronic supplementary information, we study the particular case of a Pareto-distributed loss whose severity is affected by risk reduction service. Here, under mild assumptions, $G^{\beta}(s)$ turns out to be strictly convex (with $\lim_{s\to\infty} G^{\beta}(s) = \infty$), yielding only one additional case compared to the self-protection case, namely $G^{\beta}(s)$ intersecting the level $(1+\theta)$ exactly twice.

As outlined above, we now consider a scenario of *self-protection* (Ehrlich and Becker 1972), i.e. an a.s. non-negative loss X which stems from a family of zero-inflated distributions of the form

$$F_{X,s}(x) = [(1-p(s)) + p(s)F_Y(x)]\mathbb{1}_{\{x\geq 0\}}, \tag{4}$$

where $s \mapsto p(s) \in [0,1]$ is decreasing and F_Y is the c.d.f. of an a.s. positive r.v. Y. This means that a positive loss with c.d.f. F_Y (which could describe a single loss or be a compound distribution describing a cumulative loss) occurs with a probability that can be lowered by purchasing services while the severity distribution remains untouched, relating to (R1) and (R3) above. Ansatz (4) only assumes $s \mapsto p(s)$ to be decreasing (which is natural, as increased service should decrease the loss probability). As a standard economic assumption (e.g. Courbage et al. 2013) is $s \mapsto p(s)$ being convex (decreasing marginal impact), (A2) is not necessarily implied. Therefore, we assume another sufficient condition to ensure convexity of $s \mapsto \rho_s(X)$ for distributions of the form (4), namely that both the *objective* loss probabilities $p(s)$ and the *subjective* loss probabilities $\psi(p(s))$ are decreasing in a convex way (see Bensalem et al. 2020, Lemma 3.3, and section A.3 in the electronic supplementary information).

Example 2 As ψ is concave, $s \mapsto p(s)$ must be "sufficiently" convex for the concatenation to be convex; e.g. for the common choice of distortion function $\psi(u) = u^r, r \in (0,1]$, a sufficient condition for the convexity of $\psi(p(s)) = p(s)^r$ would be for $s \mapsto p(s)$ to be *logarithmically convex* (see section A.5 in the electronic supplementary information).

Increasingness of $s \mapsto G^{\beta}(s)$ for any $\beta \in [\beta, 1]$ in the self-protection case allows a convenient expression of the sets $\mathcal{I}$ and $\mathcal{N}$.

Corollary 1 (Structure of $\mathcal{I}$ and $\mathcal{N}$ in the self-protection case, extension of Bensalem et al. (2020), *Lemma* 3.2) *There exists a constant* $\theta_0 \geq 0$ *such that*:

(1) *If* $\theta \leq \theta_0$, *then* $\mathcal{N} = \emptyset$ *and* $\mathcal{I} = [0, \infty)$.
(2) *If* $\theta > \theta_0$, *then for any* $\beta \in [\beta, 1]$, *there exists* $s_B(\theta, \beta) > 0$ *such that* $\mathcal{N} = [0, s_B(\theta, \beta))$ *and* $\mathcal{I} = [s_B(\theta, \beta), \infty)$.

In the latter case, both maps $\theta \mapsto s_B(\theta, \beta)$ and $\beta \mapsto s_B(\theta, \beta)$ are increasing.

Remark 6 (Interpretation of Corollary 1) Case (1) states that if the loading is lower than a given constant level θ_0, the buyer would purchase insurance already for the original risk (at $s = 0$) and therefore at any level s (recall that increasing s reduces the price more than the risk). Case (2), illustrated in Figure 1, corresponds to a situation where the loading is too high for the buyer to insure the original risk, but by adding a service level of at least $s_B(\theta, \beta)$ (which depends on θ as well as its relative cost β), an insurance contract with loading θ becomes acceptable for the buyer.

This directly relates to the *insurance gap* on the cyber insurance market: for the pure risk transfer ($s = 0$) policies offered with loading θ, it may not be acceptable for the buyer to insure the original risk at the price the insurer demands. To make an insurance contract possible, either θ would have to be lowered to at most a level θ_0 (move from case (2) to case (1)) or risk reduction services equivalent to a level s_B would have to be offered as part of the policy (in case (2), enable a move from $\mathcal{N}$ to $\mathcal{I}$).

Lastly, it is intuitive that if the *risk premium* or *service premium* increase, the with-insurance solution becomes relatively more expensive for the buyer, and the interval corresponding to $\mathcal{N}$ (resp. $\mathcal{I}$) becomes larger (resp. smaller).

To solve the buyer's problem, first note that $L_{1,\mathcal{N}}(s)$, resp. $L_{1,\mathcal{I}}^{\theta,\beta}(s)$, each admit a unique global minimiser on $[0, \infty)$, denoted s_N resp. $s_I(\theta, \beta)$.

Corollary 2 (Solutions of separate problems, extension of Bensalem et al. (2020, section 3.3)

1. *For any* $\beta \in [\underline{\beta}, 1]$, *there exists a positive constant* $\theta_N(\beta) > \theta_0$ *such that*

$$\theta < \theta_N(\beta) \implies argmin_{\mathcal{N}} L_{1,\mathcal{N}}(s) = s_B(\theta, \beta),$$
$$\theta \geq \theta_N(\beta) \implies argmin_{\mathcal{N}} L_{1,\mathcal{N}}(s) = s_N.$$

The map $\beta \mapsto \theta_N(\beta)$ *is decreasing.*

2. *For any* $\beta \in [\underline{\beta}, 1]$, *there exists a constant* $\theta_I(\beta)$ *such that*

$$\begin{aligned} \theta \le \theta_I(\beta) &\implies s_I(\theta,\beta) = 0, \\ \theta > \theta_I(\beta) &\implies s_I(\theta,\beta) > 0. \end{aligned}$$

In the latter case, the following hold:

(i) *For any $\beta \in [\underline{\beta}, 1]$, the map $\theta \mapsto s_I(\theta,\beta)$ is increasing.*
(ii) *For any $\theta > 0$, the map $\beta \mapsto s_I(\theta,\beta)$ is decreasing.*

Remark 7 (Interpretation of Corollary 2)

Part 1.: As the loading θ increases, the set $\mathcal{N}$ (no insurance) expands, i.e. the boundary $s_B(\theta,\beta)$ increases (shift to the right in Fig. 1). The value $\theta_N(\beta)$ is the smallest loading such that the global minimiser of $L_{1,\mathcal{N}}(s)$ lies in $\mathcal{N}$

Part 2.: For fixed service cost β, as θ increases, it becomes relatively more expensive to transfer risk, which makes it economically rational to reduce the to-be-insured risk by increasing service. Vice versa, for fixed risk loading θ, as β increases, and thus, service becomes relatively more expensive, it is economically rational to decrease the purchased amount of service.

Corollary 2 does not make a statement about the local solution on $\mathcal{I}$. As both $s_I(\theta,\beta)$ (by Corollary 2) and $s_B(\theta,\beta)$ (by Corollary 1) are non-decreasing in θ, to determine the local solution on $\mathcal{I}$ and the global solution to the minimisation of $L_1(\alpha^*(s),s)$, one has to consider all possible cases regarding the order of $s_N, s_I(\theta,\beta), s_B(\theta,\beta)$ (see sectin A.3 in the electronic supplementary information).

Corollary 3 (Global solution in the self-protection case, extension of Bensalem et al. (2020, Theorem 3.2) *For any $\beta \in [\underline{\beta}, 1]$, there exists a constant $\theta_R(\beta) \ge 0$, such that:*

(i) *If $\theta \le \theta_R(\beta)$, the global minimiser of $L_1(\alpha^*(s),s)$ is $(\alpha^*, s^*) = (1, s_I(\theta,\beta))$.*
(ii) *If $\theta > \theta_R(\beta)$, the global minimiser of $L_1(\alpha^*(s),s)$ is $(\alpha^*, s^*) = (0, s_N)$.*

Furthermore, it holds $\theta_R(\beta) \ge \theta_N(\beta)$ and the map $\beta \mapsto \theta_R(\beta)$ is non-increasing.

Fig. 1 Schematic illustration of $G^\beta(s)$ and resulting $s_B(\theta,\beta)$ for one value of θ. θ_0 is the minimum value of the loading such that $G^\beta(s)$ intersects the level $1+\theta$, resulting in $\mathcal{N}$ being non-empty

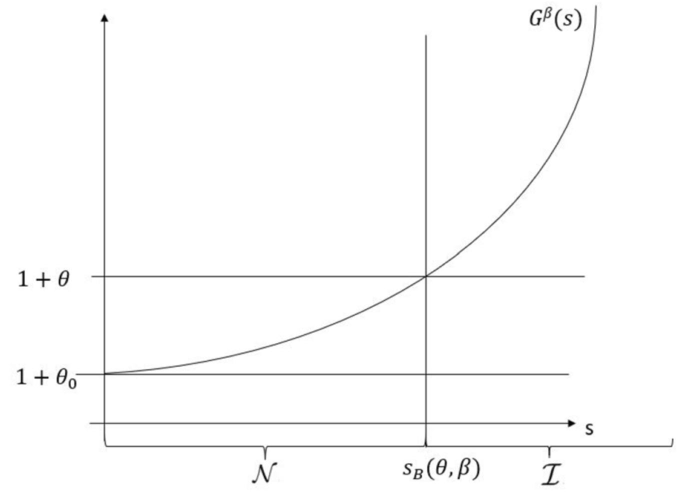

Remark 8 (Interpretation of Corollary 3) For any choice of β, there is a maximum loading $\theta_R(\beta)$ the insurance buyer is willing to accept: if it is not exceeded, he subscribes to full insurance with service level $s_I(\theta, \beta)$; else, he refrains from purchasing insurance and buys service at level s_N from an outside provider. The maximum acceptable loading decreases as the share of service cost increases, which is intuitive as the buyer accepts the contract if his total loss with insurance does not exceed his (fixed) total loss without insurance.

The relationship between risk loading and service demand is summarised in Corollary 4.

Corollary 4 (Based on Bensalem et al. 2020, Corollary 3.3) *For any $\beta \in [\underline{\beta}, 1]$, the map $\theta \mapsto s^*(\theta, \beta)$ is non-decreasing for $\theta \leq \theta_R(\beta)$ and constant (equal to s_N) for $\theta > \theta_R(\beta)$. It has a negative jump of size $s_N - s_I(\theta_R(\beta), \beta)$ at $\theta = \theta_R(\beta)$, which means that demand for risk transfer and service can be complements.*

Remark 9 (Interpretation of Corollary 4 in the cyber context) Corollary 4 is meaningful in cyber insurance: earlier game-theoretic studies concerned with the existence and efficiency of a cyber-insurance market where agents in a network invest in interdependent security measures (e.g. Lelarge and Bolot 2009; Schwartz et al. 2013; Schwartz and Sastry 2014; Shetty et al. 2010, 2010) have in many cases concluded that given the availability of cyber insurance, individuals' willingness to invest in self-protection decreases and it is, thus, generally not possible to design insurance as a means to reach socially optimal levels of investment. Corollary 4 emphasises the much more optimistic perspective that in case of self-protection, the existence of insurance can indeed lead to higher optimal levels of risk reduction at least for individual policyholders. While we do not consider negative externalities of interdependent security investments, it is reasonable to postulate that by subscribing to insurance with a high service level, policyholders inadvertently benefit other agents in their network, e.g. by reducing the risk of cyberattacks being propagated through their systems or by providing loss data the insurer can use to warn other policyholders.

Furthermore, Corollary 4 allows another understanding of the cyber insurance gap: as the optimal service demand within insurance can be higher than without insurance, for a given combination (θ, β) that an insurer demands in practice, if the service that can be offered is limited (e.g. due to technical constraints or due to limited contracts between insurers and service providers), the optimal within-insurance service level may not be attainable and the company may prefer the no-insurance solution. A way to close (or narrow) the gap would be to either decrease the premium or to increase the amount of available service within an insurance policy to make $s_I(\theta, \beta)$ attainable.

Having found the insurance buyer's optimal response to any combination (θ, β, β_o), we address the insurer's problem of choosing (θ, β) to minimise her loss over all optimal responses of the buyer.

The insurer's problem: single-contract case

Given the results of Corollary 3, (IP) reduces to a minimisation over a compact set:

$$\min_{(\theta,\beta)\in\mathcal{A}:=[0,\theta_R(\beta)]\times[\underline{\beta},1]} L_0(\theta,\beta)=\rho_{0,s_I(\theta,\beta)}(X)-(1+\theta)\mathbb{E}_{s_I(\theta,\beta)}[X]+(1-\beta)c(s_I(\theta,\beta)), \tag{5}$$

assuming that the obtainable objective value of (5) is negative. This corresponds to a choice (θ,β) yielding full risk transfer with service level $s_I(\theta,\beta)\geq 0$ as the buyer's optimal response. In case the insurer could not obtain a negative objective value in (5), she abstains from offering risk transfer by choosing $\theta>\theta_R(\beta)$ in (IP). In this case, the buyer's optimal response is $(\alpha^*,s^*)=(0,s_N(\beta_o))$, i.e. to buy service at level $s_N(\beta_o)$ outside an insurance policy.[21] Note that the special case $\beta=1$, where the insurance buyer carries the full cost of self-protection, has already been studied previously, the difference here being that the self-protection measures can be obtained cheaper within an insurance contract, increasing the maximum risk premium chargeable by the insurer.

We now state that in the self-protection case, choosing $\beta=1$ is also a solution to the more general problem (5). The steps leading to this result are outlined subsequently, proofs are postponed to section A.4 in the electronic supplementary information.

Theorem 1 (Solution of (5) *in the self-protection case*) *Let the assumptions of Lemma* 2 (*self-protection, see section A.3 in the electronic supplementary information*) *hold. Then, a solution* (θ^*,β^*) *to the minimisation problem* (5) *lies in the compact set* $\{(\theta,1)\,:\,\theta\in[0,\theta_R(1)]\}$. *This means that in the self-protection case, i.e. if service only affects the loss probability, it is always optimal for the insurer to shift the full service cost to the insured.*

Example 3 (Zero-inflated Pareto loss) The solution to (5) cannot be characterised further without more structure. Details for the special case of a zero-inflated Pareto-distributed loss are given in section A.5 of the electronic supplementary information. In this case, the insurer's loss can be shown to be monotone in θ for $\beta=1$, yielding the solution $\theta^*=\theta_R(1)$ (see Bensalem et al. 2020). Combining this with

[21] As mentioned above, one could theoretically allow the insurer to offer "service-only" contracts by solving

$$\min_{\beta\in[1,\beta_o]}(1-\beta)c(s_N(\beta)), \tag{6}$$

which certainly yields a non-positive objective value. It might be feasible to assume that the insurer would be able to offer such services cheaper than other market participants, as she might have certain service infrastructures (contracts with IT experts, warning mechanisms) in place already for her insurance clients. One might also assume that the insurer has initially solved this problem, thus, determining β_o, and the upper bound in (6) is the next-cheapest outside option. Under no circumstance would we find it realistic to allow the insurer to simultaneously compare (negative) objective values of (5) and (6) and choose the lower one. In other words, the insurer should not compare for a prospective buyer where risk transfer is profitable whether it could be more profitable to offer only services and choose a solution that discourages the buyer from buying risk transfer.

Theorem 1 means that for a Pareto-distributed loss whose occurrence probability can be lowered by risk reduction services, an optimal solution for the insurer is given by shifting the full cost of service to the insured and charging the maximum acceptable loading, i.e. $(\theta^*, \beta^*) = (\theta_R(1), 1)$.

Remark 10 Theorem 1 does not make a statement about uniqueness of the solution, as uniqueness only holds whenever the maximum attainable loading $\theta_R(\beta)$ is larger than the minimum loading $\theta_I(\beta)$ that makes pure risk transfer undesirable to the insured compared to a combination of risk reduction and risk transfer (i.e. leads to a solution $s_I(\theta, \beta) > 0$, see the proof of Corollary 2). This holds true under quite general assumptions on the function $s \mapsto c(s)$, e.g. for its right-side derivative at 0 to vanish, i.e. $c'(s)|_{s=0^+} = 0$.

We use the (implicit) definition of the maximum feasible loading for any share of service cost $\theta_R(\beta)$ from the proof of Corollary 3, given as

$$\theta_R(\beta) := \sup\left\{\theta \geq 0 : L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta)) \leq L_{1,\mathcal{N}}(s_N)\right\},$$

which is well-defined for any $\beta \in [\underline{\beta}, 1]$, as the map $\theta \mapsto L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta))$ is increasing with $L^{0,\beta}_{1,\mathcal{I}}(s_I(0,\beta)) < L_{1,\mathcal{N}}(s_N)$. Furthermore, it is shown that for any $\theta \geq 0$ (resp. $\theta > \theta_I(\beta)$), the map $\beta \mapsto L^{\theta,\beta}_{1,\mathcal{I}}(s_I)$ is non-decreasing (increasing) such that $\beta \mapsto \theta_R(\beta)$ is non-increasing (decreasing). By denoting $\underline{\theta} := \theta_R(1)$ and $\bar{\theta} := \theta_R(\underline{\beta})$, it holds $L^{\theta,\underline{\beta}}_{\mathcal{I}}(s_I(\theta,\underline{\beta})) < L_{\mathcal{N}}(s_N)$ for any $\theta \in [0, \bar{\theta}]$, such that one can likewise define for any such θ the constant

$$\beta_M(\theta) := \max\left\{\beta \in [\underline{\beta}, 1] : L^{\theta,\beta}_{\mathcal{I}}(s_I(\theta,\beta)) \leq L_{\mathcal{N}}(s_N)\right\},$$

denoting the maximum feasible share of service cost such that the contract is accepted for a given loading. The map $\theta \mapsto \beta_M(\theta)$ is by definition non-increasing on $\theta \in [0, \bar{\theta}]$. As a corollary of Lemma 2, we deduce that for $\theta \geq 0$ fixed, the insurer's loss is monotone in the share of service cost β.

Proposition 1 (Monotonicity of insurer's loss in β) *Under the conditions of Lemma* 2 (*self-protection*) *and under the necessary condition of profitability for the insurer*, *i.e. if* $L_0(\theta,\beta) < 0$, $\beta \mapsto L_0(\theta,\beta)$ *is a monotone*, *non-increasing function for any* $\theta \geq 0$.

Proposition 1 states that for any (fixed) loading θ, an optimal solution for the insurer is to choose the maximum possible service cost $\beta_M(\theta)$ acceptable to the buyer, or equivalently that the insurer has no incentive to subsidise risk reduction through a rebate on services. This implies that an optimal solution to problem (5) lies in the (compact) set $\{(\theta, \beta_M(\theta)),\ \theta \in [\underline{\theta}, \bar{\theta}]\} \cup \{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$ or equivalently $\{(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]\} \cup \{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$ (see Figure 2). The one-dimensional optimisation problem on $\{(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]\}$ can be understood as solving the insurer's trade-off between charging a higher service cost versus a

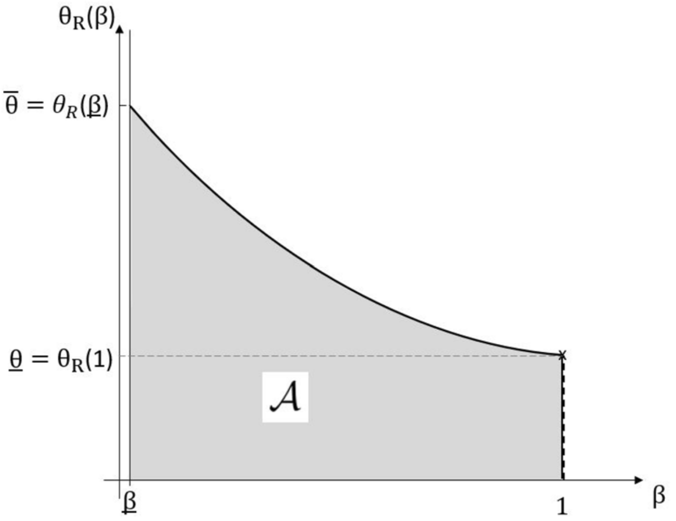

Fig. 2 Schematic illustration of the insurer's admissible set $\mathcal{A} = [0, \theta_R(\beta)] \times [\underline{\beta}, 1]$ (grey) and the set containing the optimal solution in the self-protection case. According to Proposition 1, an optimal solution must lie on the boundary $\{(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]\} \cup \{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$ (solid black line). Proposition 2 restricts the set containing an optimal solution to the set $\{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$ (dashed black line). For the special case of a Pareto-distributed loss, the optimal solution $(\theta^*, \beta^*) = (\underline{\theta}, 1)$ is marked by a cross (for details, see section A.5 in the electronic supplementary information)

higher risk loading while offering a contract the buyer will accept. The following proposition states that the insurer's loss on this set is monotone in β, leading to the statement of Theorem 1.

Proposition 2 (Monotonicity of insurer's loss in β with maximum feasible risk premium) *Under the conditions of Lemma* 2 (*self-protection*), *the map* $\beta \mapsto L_0(\theta_R(\beta), \beta)$ *is non-increasing*.

Remark 11 (Self-insurance) A central property leading to the above results for the *self-protection* case is non-decreasingness of $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$. In case of *self-insurance*, this assumption does not necessarily hold; indeed, for some standard loss distributions (e.g. Pareto, Weibull, or Log-Normal), the converse holds true, i.e. $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$ is non-increasing (see Bensalem et al. 2020). In section A.6 in the electronic supplementary information, we study the particular case of a Pareto-distributed loss whose severity is affected by risk reduction service. We find that in this self-insurance case, the insurer can indeed have an incentive to subsidise service cost (i.e. offer contracts with $\beta^* < 1$), where the optimally subsidised share $(1 - \beta^*)$ increases with the insurer's risk aversion. In particular, if the risk aversions of insurer and insurance buyer are similar (i.e. $r_0 \searrow r_1$ for the PH transform risk measure), a mutually acceptable contract may only exist if the cost is shared $(0 < \beta < 1)$. This further implies that the insurer's optimal solution, i.e. the price structure the insurance buyer is offered, may depend on his choice of risk measure, even if the initial risk assessment is equivalent.

So far, we scrutinised the interaction between the insurer and a single insurance buyer as an isolated problem. This is often reasonable, as in practice insurers usually price individual risks on a stand alone basis without taking into account the existing portfolio. However, the failure of the independence assumption between risks is one of the central challenges in cyber insurance, as cyber incidents at different firms can be dependent, e.g. due to common underlying vulnerabilities (e.g. Böhme et al. 2018; Zeller and Scherer 2022) or due to propagation for worm-type viruses. Therefore, one could argue that rather than finding price structures (θ, β) by considering problem (5) separately for each customer, the insurer should jointly optimise the risk measure for the entire portfolio against the sum of all premiums received (note that distortion risk measures are in most situations not additive for non-comonotonic risks).

In the "The insurer's problem: portfolio viewpoint" section, we illustrate that already for portfolios of two dependent losses, the results of Theorem 1 do not necessarily hold anymore, i.e. when optimising from a portfolio viewpoint, indeed the insurer can have an incentive to subsidise self-protection measures for some policyholders.

The insurer's problem: portfolio viewpoint

In the self-protection case, a central property is that for any single contract in a portfolio of n policyholders with risks X_i, $i \in \{1, \dots, n\}$, for any feasible loading θ_i, $i \in \{1, \dots, n\}$, the reduction in price for increased service outweighs the reduction in the insurer's risk measure $\rho_{0,s_i}(X_i)$, $i \in \{1, \dots, n\}$ for each single risk, i.e.

$$\underbrace{(1+\theta_i)\frac{\partial \mathbb{E}_{s_i}[X_i]}{\partial s_i}}_{\text{sensitivity of premium for } X_i} < \underbrace{\frac{\partial \rho_{0,s_i}(X_i)}{\partial s_i}}_{\text{sensitivity of risk measure for } X_i} \leq 0, \quad i \in \{1, \dots, n\}.$$

However, ordering of the relevant sensitivities is not necessarily preserved in a portfolio context, i.e. when adding a new policyholder to an existing portfolio, the reduction of the overall *portfolio risk measure* $\rho_{0,\mathbf{s}}(X)$ may outweigh the price reduction of the additional contract, i.e. for some $i \in \{1, \dots, n\}$:

$$\frac{\partial \rho_{0,\mathbf{s}}(X)}{\partial s_i} < (1+\theta_i)\frac{\partial \mathbb{E}_{s_i}[X_i]}{\partial s_i} < \frac{\partial \rho_{0,s_i}(X_i)}{\partial s_i} \leq 0, \tag{7}$$

where $\mathbf{s} := (s_1, \dots, s_n)$ and $X = \sum_{i=1}^n X_i$ is the aggregated loss. This may imply a situation where the insurer has an economic incentive to subsidise risk reduction for some policyholders in the self-protection case, as we will now analyse in a toy example of two policyholders with dependence mechanisms representative for cyber risk: (directed) loss propagation, common cyber events, and copula approaches. While these bivariate examples will already be sufficient to work out the structural difference to the univariate case, we provide one exemplary extension to a general multivariate setting in section A.7.3 of the electronic supplementary information.

(Directed) loss propagation

A popular way of modelling dependencies between cyber losses is to consider a model of epidemic spreading in an underlying network, i.e. a directed or undirected graph whose nodes are interpreted as companies (or machines) and whose edges are interpreted as connections between these companies (or machines) through which a state of "infectiousness" can be passed on. These models, often originating from mathematical biology, have been extensively studied in the cyber context over the last few years, see, e.g. Fahrenwaldt et al. (2018), Xu and Hua (2019), Xu et al. (2015) or the surveys Marotta et al. (2017), and Kerstin Awiszus et al. (2022). Interpretations of such models are worm-type viruses spreading between connected machines or a state of business interruption propagating through a supply chain.

Example 4 (Bivariate model with one directed edge) For illustration purposes, we consider a portfolio of two firms with one directed edge between them and we understand the "infected" state as a loss occurrence, i.e. assume a loss occurrence in firm 1 can cause a loss in firm 2 with probability $q \in [0,1]$, but not vice versa.[22] If a loss occurs, the loss sizes are deterministic; w.l.o.g. $0 < L_1 \leq L_2 < \infty$. We assume that the events of the occurrence of a loss in firm 1, its propagation, and the occurrence of a non-propagated loss in firm 2 are independent. This implies that, depending on the chosen service levels $s_i,\ i \in \{1,2\}$, the loss r.v.s $X_i,\ i \in \{1,2\}$, take the values

$$X_1 = \begin{cases} 0 & \text{w.p. } 1 - p_1(s_1), \\ L_1 & \text{w.p. } p_1(s_1), \end{cases} \quad X_2 = \begin{cases} 0 & \text{w.p. } 1 - (p_2(s_2) + qp_1(s_1)(1 - p_2(s_2))), \\ L_2 & \text{w.p. } p_2(s_2) + qp_1(s_1)(1 - p_2(s_2)), \end{cases}$$

where $s \mapsto p_i(s)$ are continuous, non-increasing functions with $\lim_{s\to\infty} p_i(s) > 0$ for $i \in \{1,2\}$. Let $X := X_1 + X_2$ denote the portfolio loss, such that the insurer's portfolio risk measure, using $\psi(u) = u^{r_0},\ r_0 \in (0,1]$, is given by (see section A.7.1 in the electronic supplementary information):

$$\rho_{0,\mathbf{s}}(X) = L_1[(p_1 + p_2 - p_1p_2)^{r_0} + (p_1q + p_1p_2 - p_1p_2q)^{r_0}] + (L_2 - L_1)(p_2 + p_1q - p_1p_2q)^{r_0},$$

where the dependence on $s_i,\ i \in \{1,2\}$, is suppressed for notational convenience and $\mathbf{s} := (s_1, s_2)$.

Figure 3 illustrates that (7) may hold in the above example, which indicates that the insurer can have a financial incentive to subsidise service.

Remark 12 (Insurer's problem: individual optimisation) If the insurer evaluates the two contracts individually, she solves separately

[22] The cited works typically use two processes, one to model the state of infectiousness among nodes in the graph and another one for loss occurrences among "infected" nodes; we regard this additional complexity as unnecessary for the present example.

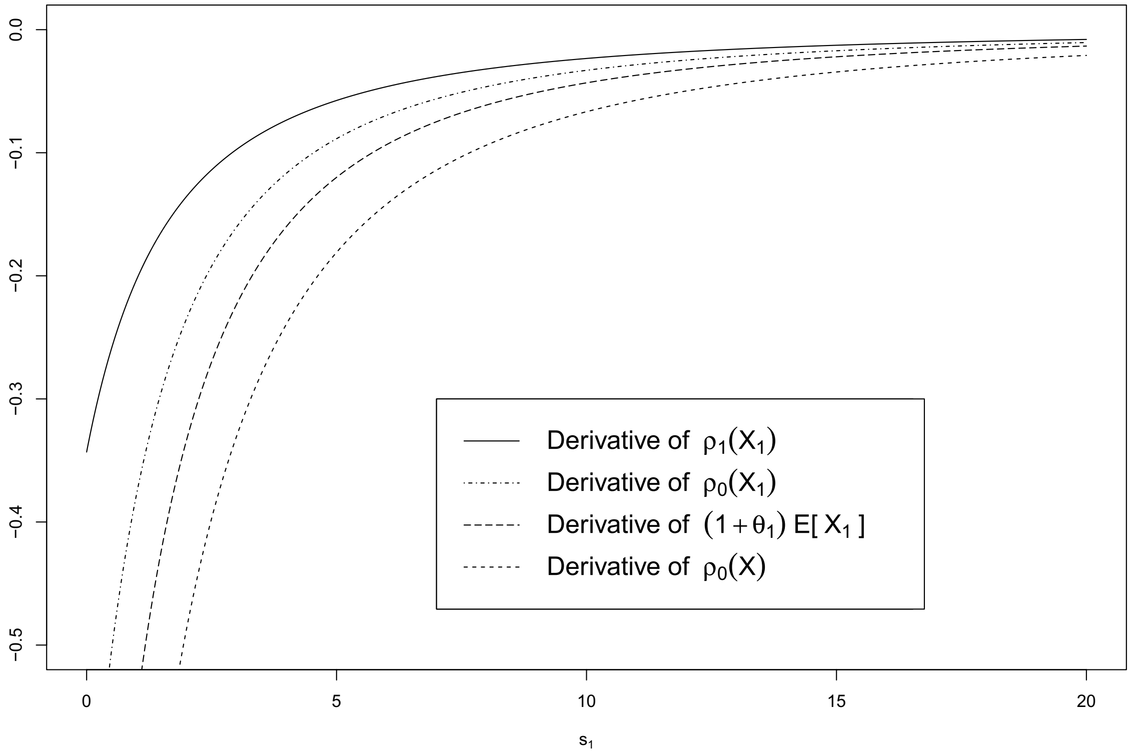

Fig. 3 Comparison of derivatives with respect to s_1 of single-contract and portfolio risk measures as well as the price of insurance (at a feasible loading $\theta_1 = 0.35$). Note that Equation (7) holds: The decrease in price outweighs the decrease in both single-contract risk measures, but is outweighed by the reduction in the insurer's portfolio risk measure. The parameters for this example are chosen as $r_0 = 0.8,\ r_1 = r_2 = 0.3,\ L_1 = 5,\ L_2 = 10,\ p_1(s_1) = \frac{1}{a+s_1} + b = \frac{1}{2.5+s_1} + 0.2,\ p_2 = 0.3,\ q = 0.8$

$$\min_{(\theta_i,\beta_i)\in[0,\theta_{R,i}(\beta_i)]\times[\underline{\beta},1]} L^{\text{ind}}_{0,i}(\theta_i,\beta_i) = \rho_{0,s_{Ii}(\theta_i,\beta_i)}(X_i) - (1+\theta_i)\mathbb{E}_{s_{Ii}(\theta_i,\beta_i)}[X_i] + (1-\beta_i)c(s_{I,i}(\theta_i,\beta_i)),\ i \in \{1,2\}, \tag{8}$$

where the superscript 'ind' denotes *individual* contract pricing.

Remark 13 By very similar calculations as for the Pareto case, one can show that for a loss of deterministic severity, $\theta \mapsto L_{0,i}(\theta, 1)$ is monotone non-increasing, such that the insurer's optimal solution to the minimisation problems (8) is $(\theta_i^*, \beta_i^*) = (\theta_{R,i}(1), 1),\ i \in \{1,2\}$, i.e. to shift the full cost of service to the buyers and charge the maximum feasible loading, respectively.

We now consider her optimisation problem from a portfolio viewpoint in a two-contract set-up, where, interestingly, it has to be distinguished whether the contracts with the buyers are closed sequentially or simultaneously. Let us commence by assuming that the two contracts are closed sequentially and firm 2 is insured first.

Example 5 (Interpretation of sequential contract closure) Sequential contract closure could be interpreted as a situation where for a prospective policyholder, a loss could be caused by an occurrence at another firm (e.g. a supplier) outside the insurer's portfolio, but insuring the other firm is not feasible (yet).

Remark 14 (Insurer's problem: sequential optimisation, first policy) The results for firm 2, being insured first, are analogous to the single-contract case: In her initial risk assessment, assume the insurer correctly assesses the loss probability (given service level s_2) as

$$\mathbb{P}_{\mathbf{s}}(X_2 = L_2) = p_2(s_2) + qp_1(s_1)\big(1 - p_2(s_2)\big), \tag{9}$$

which depends (due to loss propagation) on the unknown loss probability of firm 1.[23] For this study, we assume that firm 1 has not subscribed to insurance yet, but has solved the minimisation problem for the no-insurance case correctly, such that in Eq. (9) we set $s_1 = s_{N1}$. As remarked above, we know that the solution to the insurer's problem (8) for $i = 2$ is given by $(\theta_2^*, \beta_2^*) = (\theta_{R,2}(1), 1)$ and given (9), we can proceed analogously to Sect. 3 to deduce firm 2's optimal service level without insurance s_{N2} and $s_{I2}(\theta_2^*, \beta_2^*)$ within insurance.

The striking observation is as follows: By incentivising a higher service level in a subsequent contract with firm 1, the insurer not only improves the to-be-insured risk in that contract, but also the already priced risk in the existing contract with firm 2, as the probability for a propagated loss decreases.[24]

Remark 15 (Insurer's problem: sequential optimisation, second policy) If the insurer prices each contract as if the risks were independent (or the propagation potential is undetected), she would solve (8) for $i = 1$ yielding $(\theta_1^*, \beta_1^*) = \big(\theta_{R,1}(1), 1\big)$. However, if she correctly takes the effect on the portfolio risk into account, to find (θ_1^*, β_1^*) she instead considers the problem

$$\begin{aligned}\min_{(\theta_1,\beta_1)\in[0,\theta_{R,1}(\beta_1)]\times[\underline{\beta},1]} L_{0,1}^{\text{seq}}(\theta_1, \beta_1) =& \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_{R,2}(1),1)}(X) \\ &- (1+\theta_1)\mathbb{E}_{s_{I1}(\theta_1,\beta_1)}(X_1) - (1+\theta_2)\mathbb{E}_{s_{I2}(\theta_{2,R}(1),1)}(X_2) \\ &+ (1-\beta_1)c(s_{I1}(\theta_1,\beta_1)) + (1-\beta_2)c(s_{I2}(\theta_{2,R}(1),1)),\end{aligned} \tag{10}$$

where the superscript 'seq' denotes *sequential* contract closure and $X = X_1 + X_2$.[25]

Remark 16 Sequential contract closure in the reverse order can be studied analogously. It is, however, obvious from the set-up of directed loss propagation that the insurer has no additional incentive to subsidise service for firm 2, independently of whether firm 1 is part of the portfolio, i.e. this analysis would not yield different results from the single-contract case and is, thus, omitted.

[23] Note that (9) implies that by buying service from the insurer, firm 2 can reduce the probability of a non-propagated loss only.

[24] This example is somewhat related to the question studied in Khalili et al. (2019) about jointly underwriting a service provider and its customers as interdependent insurance customers.

[25] Note that the terms corresponding to contract 2 are not adjusted at this point and therefore will not appear in the minimisation.

We now assume that both contracts are priced simultaneously.

Example 6 (Interpretation of simultaneous contract closure) In practice, simultaneous contract closure could be interpreted as two firms jointly inquiring about insurance (e.g. companies along a supply chain or parent company and subsidiary) or the insurer approaching both before the first contract is closed.

Remark 17 (Insurer's problem: simultaneous optimisation) If the insurer offers both contracts simultaneously, she considers the four-dimensional problem

$$\begin{aligned}\min_{(\theta_1,\beta_1,\theta_2,\beta_2)\in\mathcal{A}} L_0^{\mathrm{sim}}(\theta_1,\beta_1,\theta_2,\beta_2) = {} & \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)}(X) \\ & -(1+\theta_1)\mathbb{E}_{s_{I1}(\theta_1,\beta_1)}[X_1]-(1+\theta_2)\mathbb{E}_{s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)}[X_2] \\ & +(1-\beta_1)c(s_{I1}(\theta_1,\beta_1))+(1-\beta_2)c(s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)),\end{aligned} \tag{11}$$

where the superscript 'sim' denotes *simultaneous* contract closure, $X = X_1 + X_2$, and $\mathcal{A} := [0, \theta_{R,1}(\beta_1)] \times [\underline{\beta}, 1] \times [0, \theta_{R,2}(\beta_2)] \times [\underline{\beta}, 1]$ is the admissible set for this problem.

The results of numerically solving the above optimisation problems are given in Fig. 4 for the propagation probability $q \in [0, 1]$, which in this set-up governs the dependence between the risks.[26]

Remark 18 (Interpretation of results for directed loss propagation)

- **Panel 4(a)** depicts the optimal pricing parameters (θ_1^*, β_1^*) of the contract offered to firm 1 (the "source of propagation"). If the contract with firm 2 is priced first, the insurer may subsidise service (i.e. choose $\beta^* < 1$) in the subsequent contract with firm 1, as this reduces the insured risk in contract 2 (without having to adjust the premium of firm 2). This subsidy $(1 - \beta^*)$, as well as the loading θ_1^*, increase with the dependence between the risks. The same effect occurs, but to a smaller extent, if the contracts are priced simultaneously. This is caused by the fact that by subsidising service for firm 1, the insured risk in firm 2 is reduced, but this now has to be reflected in a decreased chargeable premium for that contract. Therefore, the incentive to subsidise service for firm 1 is smaller relative to the case where the price of contract 2 is fixed first.
- **Panel 4(b)** depicts the optimal parameters (θ_2^*, β_2^*) of the contract offered to firm 2. As the service level of firm 2 has no additional effect on firm 1, the insurer's problem for firm 2 is always analogous to the single-contract case, and thus, service cost is never subsidised $(\beta^* = 1)$. However, the risk loading depends on the

[26] The calculation of the gradients, used in the numerical optimisation routine, is detailed in section A.7.1 of the electronic supplementary information.

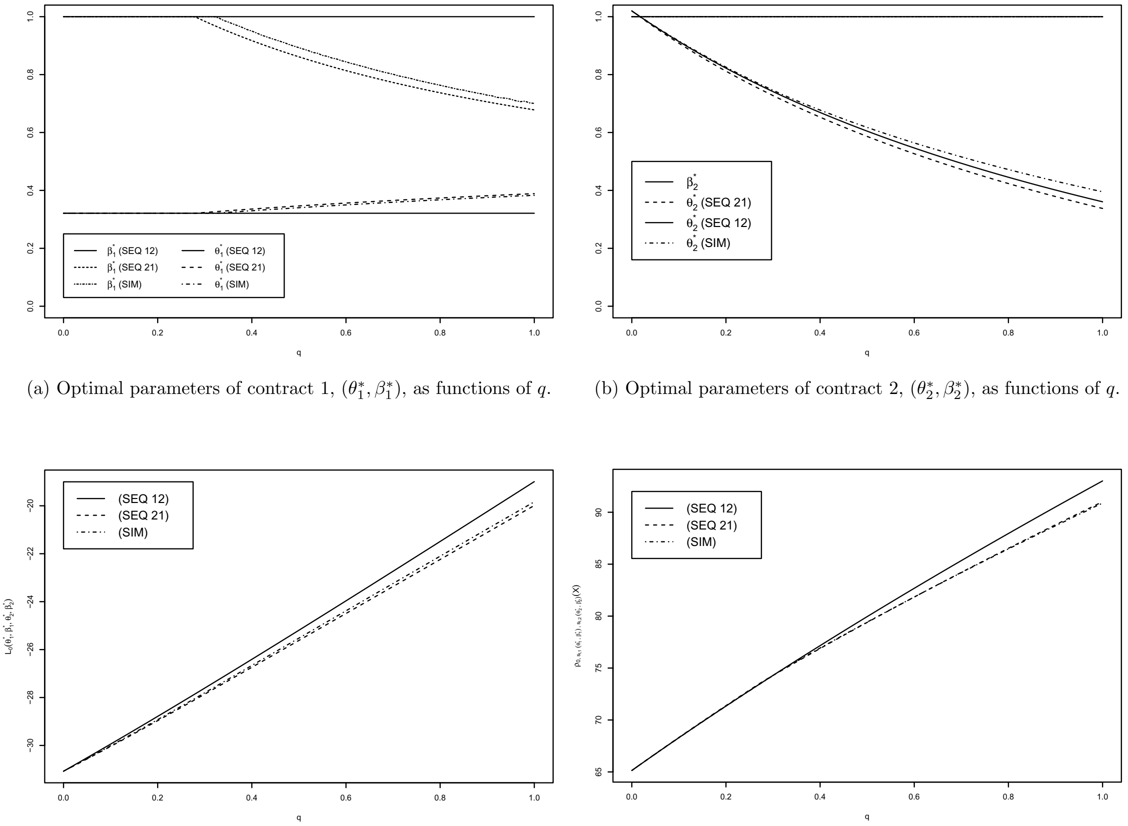

(a) Optimal parameters of contract 1, (θ_1^*, β_1^*), as functions of q.

(b) Optimal parameters of contract 2, (θ_2^*, β_2^*), as functions of q.

(c) Insurer's optimal attainable loss $L_0(\theta_1^*, \beta_1^*, \theta_2^*, \beta_2^*)$.

(d) Insurer's optimal portfolio risk $\rho_{0,s_{I1}(\theta_1^*,\beta_1^*),s_{I2}(\theta_2^*,\beta_2^*)}(X)$.

Fig. 4 Aspects of the insurer's solution in the portfolio case with directed loss propagation. The parameters for this example are: loss sizes $L_1 = 50, L_2 = 100$, loss probability parameters $a_1 = a_2 = 2.5, b_1 = b_2 = 0.2$, risk aversion $r_0 = 0.8, r_1 = 0.7, r_2 = 0.3$, cost parameters $\eta = 0.5, \gamma = 2, \beta_o = 1.1, q \in [0, 1]$

loss probability $\mathbb{P}_{\mathbf{s}}(X_2 = L_2)$, which differs between the cases as it depends on s_1^* and therefore on whether firm 1 is insured already (and under which parameters).

- **Panel 4(c)** depicts the insurer's optimally attainable negative loss (gain) $L_0(\theta_1^*, \beta_1^*, \theta_2^*, \beta_2^*)$, which decreases with increasing dependence between the risks, while the additional gain from pricing contracts "correctly", i.e. using the portfolio risk measure, increases with the dependence. Analogous observations hold for the insurer's portfolio risk, see Panel 4(d).

Cyber events at multiple 'targets'

Another way to understand dependence between cyber losses is to consider the presence of common (*systemic*) vulnerabilities which allow cyber threats to affect multiple companies simultaneously (see, e.g. Böhme et al. 2018; Zeller and Scherer 2022). Realistic examples for systemic events causing incidents in multiple firms are the accidental outage or the malicious exploitation of a vulnerability in commonly used software or operating systems, leading to, e.g. data breaches or fraudulent activity (e.g. ransomware claims).[27]

[27] One recent example were the multiple, sometimes effectively simultaneous attacks on exchange servers via the so-called *ProxyShell* exploit during 2021, see, e.g. Born (2021).

Remark 19 (Buyer's vs. insurer's perspective on common events) In this setting, each company faces incidents from systemic events as well as *idiosyncratic* incidents occurring independently from other firms, e.g. the loss or theft of hardware or negligent employee behaviour leading to involuntary data disclosure or business interruption. From the viewpoint of each company (insurance buyer), both types of incidents are indistinguishable in the sense that they aggregate to one loss arrival process, i.e. the company simply monitors if a loss occurs (disregarding its source) without knowing (or caring) if others may be simultaneously affected. From the insurer's portfolio viewpoint, however, the two types of incidents are viewed differently: incidents from systemic events are particularly worrisome as they entail accumulation risk, whereas idiosyncratic incidents are "desirable" in the sense that they constitute (if correctly priced) the basis of the insurance business and can be "diversified away" in a large portfolio.

Example 7 (Bivariate model with common events) Consider as model for the risks X_1 and X_2:

$$X_1 = L_1 \mathbf{1}_{\{\min\{E_1, E_{12}\} \leq T\}}, \; X_2 = L_2 \mathbf{1}_{\{\min\{E_2, E_{12}\} \leq T\}},$$

with $E_1 \sim Exp(\lambda_1)$, $E_2 \sim Exp(\lambda_2)$, and $E_{12} \sim Exp(\lambda_{12})$ independent with $\lambda_1, \lambda_2, \lambda_{12} \geq 0$, *s.t.* $\lambda_i + \lambda_{12} > 0$, $i \in \{1, 2\}$, and w.l.o.g. $0 < L_1 \leq L_2 < \infty$. E_1 and E_2 model the arrival times of an idiosyncratic incident to firm 1 and 2, respectively, whereas E_{12} models the arrival time of a common event causing simultaneous incidents in both firms, with deterministic loss sizes L_1 and L_2, respectively. Let T denote the time horizon of the policy under consideration (w.l.o.g. $T = 1$ in what follows) and let

$$\lambda_I := \lambda_1 + \lambda_{12}, \; \lambda_{II} := \lambda_2 + \lambda_{12},$$

denote the overall marginal arrival rates of incidents to firms 1 and 2, respectively.[28] It follows that the buyers' risk measure and expected loss are given by

$$\rho_1(X_1) = L_1(1 - e^{-\lambda_I})^{r_1}, \; \mathbb{E}[X_1] = L_1(1 - e^{-\lambda_I}),$$
$$\rho_2(X_2) = L_2(1 - e^{-\lambda_{II}})^{r_2}, \; \mathbb{E}[X_2] = L_2(1 - e^{-\lambda_{II}}),$$

while the insurer's portfolio risk measure is given by (see section A.7.2 in the electronic supplementary information)

$$\rho_0(X) = L_1[(1 - y_{00})^{r_0} + (1 - (y_{00} + y_{10} + y_{01}))^{r_0}] + (L_2 - L_1)(1 - (y_{00} + y_{10}))^{r_0},$$

where $y_{00} := e^{-(\lambda_1 + \lambda_2 + \lambda_{12})}$, $y_{10} := (1 - e^{-\lambda_1})e^{-(\lambda_2 + \lambda_{12})}$, $y_{01} := (1 - e^{-\lambda_2})e^{-(\lambda_1 + \lambda_{12})}$ are the probabilities of none (subscript $_{00}$) or exactly one (subscripts $_{10}$ and $_{01}$) of the companies experiencing a loss.[29]

[28] This corresponds to the seminal Marshall–Olkin shock model, see Marshall and Olkin (1967).

[29] In this case X_1 and X_2 are comonotone iff $\lambda_1 = \lambda_2 = 0$, implying $y_{00} = e^{-\lambda_{12}}$, $y_{10} = y_{01} = 0$ such that the risk measure is additive (a well-known general result for DRM): $\rho(X_1) + \rho(X_2) = \rho(X) = (L_1 + L_2)(1 - e^{-\lambda_{12}})^{r_0}$.

Remark 20 (Interpretation: Self-protection by prevention of systemic events) We now consider the effect of self-protection services which can be distinguished into different categories described in Table 4. In the following, we scrutinise one possible type of effect we regard as particularly interesting in the cyber context, namely the prevention of systemic events: as the existence of common vulnerabilities (e.g. use of the same software) is regarded as the source of dependence between losses, it is firstly crucial for a cyber insurer to identify such common factors among policyholders and offer services which prevent the manifestation of a loss from a systemic event for the policyholder himself (e.g. timely patch management for standard software). Second, it is in the insurer's interest to use knowledge about an incident (or so-called *near miss*, i.e. a threat that did not lead to an incident due to adequate controls) at one insured company to immediately warn other policyholders about the imminent threat and, thus, hopefully increase the chance of averting a loss manifestation for them. Thus, the total portfolio loss in case of a systemic event could be reduced or, if **all** policyholders are warned on time, the manifestation of the systemic event could even be prevented.[30]

Remark 21 (Insurer's problem: sequential optimisation, first policy) Assume again sequential contract closure, where w.l.o.g. the contract with firm 2 is closed first and its chosen service level affects the rate λ_{II} via a decreasing map $s_2 \mapsto \lambda_{II}(s_2)$.[31] Recall that by Lemma 3 (see section A.3 in the electronic supplementary information) a sufficient condition for convexity of the insurance buyer's optimisation problem is to choose the map $s_2 \mapsto \lambda_{II}(s_2)$ in such a way that the subjective loss probability $s_2 \mapsto \psi_2\big(\mathbb{P}_{\mathbf{s}}(X_2 = L_2)\big) = (1 - e^{-\lambda_{II}(s_2)})^{r_2}$ is convex. For simplicity, we choose analogously to above (however, for the rate, not the loss probability directly)

$$\lambda_{II}(s_2) = \frac{1}{s_2 + a_2} + b_2,$$

with $a_2, b_2 > 0$ such that the above convexity condition is fulfilled.

With the contract closure of firm 2, the insurer solves the single-contract problem (8) for $i = 2$, resulting in $(\theta_2^*, \beta_2^*) = \big(\theta_{R,2}(1), 1\big)$ and within-insurance service level $s_2^* = s_{I2}\big(\theta_{R,2}(1), 1\big)$ determining the loss probability of firm 2 via the rate $\lambda_{II}(s_2^*)$.

[30] Our model implicitly equates incident arrival times (e.g. $Z_1 := \min\{E_1, E_{12}\}$) with loss occurrence times, which would not allow time for a warning mechanism as all losses occur instantly and simultaneously. In reality, however, the discovery and exploitation of the same vulnerability in different firms can be delayed over time, see again, e.g. the *ProxyShell* exploit case (Born 2021). As we do not take into account discounting over the policy year and therefore do not need to explicitly model a delayed loss occurrence time after the incident arrival time, we assume the warning mechanism to directly prevent the incident arrival.

[31] In this sequential set-up, there is no distinction between *idiosyncratic* incidents and incidents from *systemic* events yet, as firm 1 is not yet part of the portfolio; in other words, the overall rate $\lambda_{II} = \lambda_2 + \lambda_{12}$ can be observed, but it is not yet distinguished between λ_2 and λ_{12}.

Table 4 Different effects of risk mitigation service if systemic events are the source of loss dependence

Scenario	Model	Effect on marginals and dependence			Interpretation
		$s_1 \mapsto \lambda_I(\mathbf{s})$	$s_1 \mapsto \lambda_{II}(\mathbf{s})$	$s_1 \mapsto 1 - \frac{\lambda_{1(2)}(\mathbf{s})}{\lambda_{I(II)}(\mathbf{s})}$	
Prevention of idiosyncratic incidents	s_1 affects risk X_1 via decreasing map $s_1 \mapsto \lambda_1(s_1)$	Decreases	Constant	Increases	Service prevents idiosyncratic incidents, e.g. continuous monitoring and improvement of password and access control management to impede unauthorised access to confidential data or processes.
Prevention of systemic events	s_1 affects portfolio risk X via decreasing map $s_1 \mapsto \lambda_{12}(s_1)$	Decreases	Decreases	Decreases	Service prevents the manifestation of a loss from a systemic event for the policyholder and allows the insurer to prevent a potential loss from the same source in other companies in the portfolio. An example is timely patch management for all common software where additionally all *near misses* are immediately reported to and analysed by the insurer (or her service provider), allowing them to identify current threats and warn other firms.

Table 4 (continued)

Scenario	Model	Effect on marginals and dependence			Interpretation
		$s_1 \mapsto \lambda_I(\mathbf{s})$	$s_1 \mapsto \lambda_{II}(\mathbf{s})$	$s_1 \mapsto 1 - \frac{\lambda_{1(2)}(\mathbf{s})}{\lambda_{I(II)}(\mathbf{s})}$	
Transformation of systemic events to idiosyncratic incidents	s_1 affects portfolio risk X via decreasing map $s_1 \mapsto \lambda_{12}(s_1)$, such that for fixed s_2, $s_1 \mapsto \lambda_{II}(\mathbf{s})$ is constant	Decreases	Constant	Decreases	Service makes firm 1 less frequently affected by incidents from systemic events, e.g. improved patch management for common software or usage of different operating systems or cloud service providers, but the risk for other firms in the portfolio is not improved (i.e. all events that would have affected both firms jointly initially now affect firm 2 alone, such that the risk X_2 does not change).[a]
	s_1 affects portfolio risk X via decreasing map $s_1 \mapsto \lambda_{12}(s_1)$, such that $s_1 \mapsto \lambda_I(\mathbf{s})$ is constant	Constant	Decreases	Decreases	Service in the contract of firm 1 does not prevent a loss in firm 1, but allows the insurer to warn others. It is obvious that firm 1 has no economic incentive to purchase such service (i.e. any $\beta_1 > 0$ yields $s_{I1}(\theta_1, \beta_1) = 0$ for any $\theta > 0$), such that in this case one would have to allow $\beta_1 = 0$ and reformulate the insurer's optimisation in terms of θ_1 and s_1, where s_1 is the amount of service the insurer would optimally include for free as part of the policy to optimise her portfolio risk.

[a] Böhme (2005) analyses the similar idea of premium discrimination between users of a dominant and an alternative platform (e.g. representing an operating system) to estimate the extent to which insurance premiums can motivate “ecosystem diversification” and counterbalance market processes that converge to a “monoculture” of installed systems

Remark 22 (Insurer's problem: sequential optimisation, second policy) At subsequent contract offering to firm 1, we assume that the service level of firm 1 influences the rate $\lambda_I(s_1)$ via a decreasing map

$$s_1 \mapsto \lambda_{12}(s_1) = \frac{1}{s_1 + a_{12}} + b_{12},$$

with $a_{12}, b_{12} > 0$ such that $s_1 \mapsto \psi_1\big(\mathbb{P}_{\mathbf{s}}(X_1 = L_1)\big)$ is convex and it must hold $\lambda_{12}(s_1) \leq \lambda_{II}(s_2^*)$ for any $s_1 \geq 0$. The marginal rates for both firms are then given by (now the incidents can be classified as idiosyncratic or systemic)

$$\lambda_I(s_1) = \lambda_1 + \lambda_{12}(s_1),$$
$$\lambda_{II}(s_1, s_2^*) = \lambda_2(s_2^*) + \lambda_{12}(s_1),$$

for some constant $\lambda_1 > 0$, implying that the choice of s_1 affects the marginal distributions of both risks as well as the dependence between them, e.g. expressed by $s_1 \mapsto 1 - \frac{\lambda_1}{\lambda_I(s_1)}$.[32] Therefore, when offering a contract to firm 1, the insurer should again consider problem (10) to correctly take the dependence into account, as opposed to solving (8) for $i = 1$.

Remark 23 (Results for prevention of systemic events) Numerical results of solving (10) are given in Fig. 5 for varying degree of dependence between the two risks.[33] We observe that if the contract of firm 1 is priced using (10), it can be optimal for the insurer to choose $\beta_1^* < 1$, leading to an increased risk loading, an increased optimal service level s_{I1} within the insurance policy, a decreased loss probability for both policyholders, and an increased gain and decreased portfolio risk for the insurer. These effects increase with the dependence between the two risks.[34]

Copula approaches

Copula approaches have become a widely popular method to assess and describe dependence between random variables, as they allow the decomposition of a multivariate distribution function (c.d.f.) F of a random vector $(X_1, \dots, X_d)$ into marginal c.d.f.s $F_1, \dots, F_d$ and an object representing the *dependence structure*, called copula C, which itself is a multivariate c.d.f. with standardized uniform marginals (see section A.2 in the electronic supplementary information). In empirical research on cyber-risk modelling, one starts with observations of cyber losses that

[32] Note that in this set-up, neither independence nor comonotonicity can be reached, as $b_{12} > 0$ and $\lambda_1 > 0$, respectively.

[33] The gradients used for the numerical optimisation are given in section A.7.2 in the electronic supplementary information. Due to the symmetrical set-up of the dependence, we do not consider the reverse order of contract closures.

[34] Note that contrary to the last example, the x-axis does not start at $\lambda_{12}(0) = 0$ representing (initial) independence, resulting in $\beta_1^* < 1$ for the whole depicted range $\lambda_{12}(0) \in \{0.15, 2\}$.

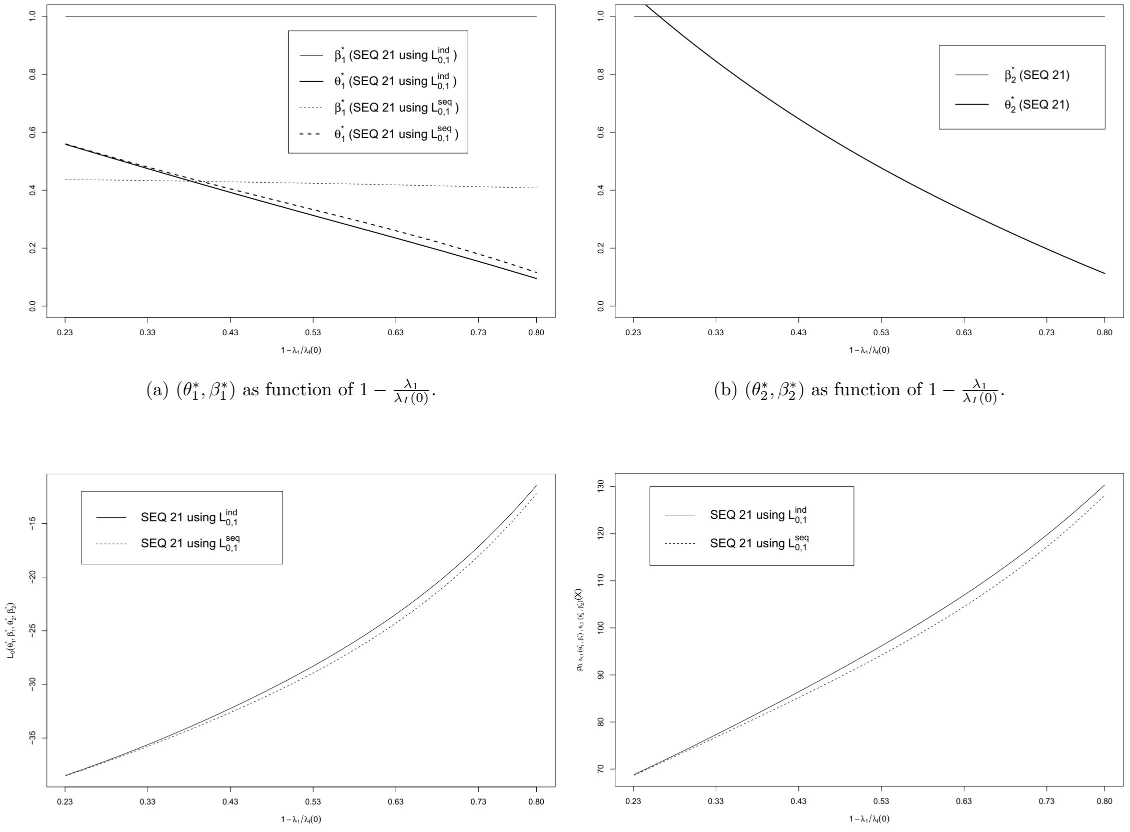

(a) (θ_1^*, β_1^*) as function of $1 - \frac{\lambda_1}{\lambda_I(0)}$.

(b) (θ_2^*, β_2^*) as function of $1 - \frac{\lambda_1}{\lambda_I(0)}$.

(c) Insurer's optimal loss $L_0(\theta_1^*, \beta_1^*, \theta_2^*, \beta_2^*)$.

(d) Insurer's optimal portfolio risk $\rho_{0,s_{I1}(\theta_1^*,\beta_1^*),s_{I2}(\theta_2^*,\beta_2^*)}(X)$.

Fig. 5 Aspects of the insurer's solution in the portfolio case with common cyber events. The parameters for this example are: $L_1 = 50, L_2 = 100$, $\lambda_1 = 0.5$, $r_0 = 0.8, r_1 = 0.4, r_2 = 0.3$, $\eta = 0.5, \gamma = 2, \beta_o = 1.1$, with $\lambda_{12}(0) = \frac{1}{a_{12}} + b_{12} \in [0.15, 2]$ and $\frac{\lim_{s\to\infty} \lambda_{12}(s)}{\lambda_{12}(0)} = \frac{1}{2}$ for any starting value

are conjectured not to be independent. As the main goal of many empirical studies is the description and analysis of the observed data, *bottom-up* approaches that seek to mimic the mechanism underlying the dependence between cyber losses may not be available for a statistical investigation, yet. Rather, a *top-down* approach of analysing the multivariate observations by fitting (parametrically or non-parametrically) univariate distributions to the marginals and by choosing a flexible parametric copula family and fitting its parameter(s) to the observed data, is often preferred (due to numerical tractability).

In the cyber context, e.g. Eling and Jung (2018) study the cross-sectional dependence of data breach losses (cross-industry and cross-breach type) using a *Gaussian* copula, among others. Previously, Böhme and Kataria (2006) and Herath and Herath (2011) proposed models for cyber risk using the *t-copula* and the *Archimedean copula family (Clayton and Gumbel)*, respectively. More recently, Peng et al. (2018) studied the multivariate dependence exhibited by real-world cyber attack data using a Copula-GARCH model with *vine copulas*.

Example 8 (Bivariate Gumbel copula) An example akin to the ones above would be for the bivariate case $(X_1, X_2) \sim F_s$ with

$$F_{\mathbf{s}}(x_1, x_2) = C_{\theta(\mathbf{s})}\big(F_{1,s_1}(x_1), F_{2,s_2}(x_2)\big), \quad x_1, x_2 \in \mathbb{R},$$

where F_{i,s_i} are the marginal c.d.f.s of the single risks depending on the chosen service levels s_i (for example, *zero-inflated Pareto distributions* as considered in Appendix 7.5 in elctronic supplementary information) and $C_{\theta(\mathbf{s})}(u, v)$ is the bivariate Gumbel copula (see Gumbel 1960)

$$C_{\theta(\mathbf{s})}(u, v) = \exp\Big[-\big((-\ln(u))^{\theta(\mathbf{s})} + (-\ln(v))^{\theta(\mathbf{s})}\big)^{1/\theta(\mathbf{s})}\Big], \quad \theta(\mathbf{s}) \in [1, \infty), \; u, v \in [0, 1],$$

which seems a suitable choice in the cyber-risk context as it allows for capturing *upper tail dependence* and is the only member of the Archimedean family which is also an *extreme-value copula*.[35] The dependence is governed by the parameter $\theta(\mathbf{s}) \in [1, \infty)$, ranging between the *independence* copula for $\theta(\mathbf{s}) = 1$ and perfect positive dependence (i.e. converging to the *comonotonicity* copula) for $\theta(\mathbf{s}) \to \infty$.[36]

Remark 24 (Effects of service on portfolio risk in the copula setting) Again, different assumptions about how the chosen service levels $\mathbf{s} = (s_1, s_2)$ of insurance buyers influence the (joint) portfolio risk can be postulated:

- If service only influences the marginal distribution of the insured risk, i.e. via $s_i \mapsto F_{i,s_i}$, $i \in \{1, 2\}$, inducing a decreasing order in the sense of the "Model set-up and assumptions" section, the analysis does not differ from the univariate case. For examples in the cyber context, see the first row of Table 4.
- If service only affects the dependence between the risks via a (in some suitable (partial) ordering decreasing) map $\mathbf{s} \mapsto \theta(\mathbf{s})$ without altering the marginals, it is obvious that no insurance buyer would have an economic incentive to purchase such service (compare the last case in Table 4) and another (interesting!) question would arise, namely, how much the insurer should optimally spend on giving away service (as a free addition to risk transfer) to favourably (in her risk measure) alter the dependence structure of her portfolio.
- If service affects both the marginal distribution(s) and the dependence structure, an example where both parties agree to share the cost of service could be constructed. For interpretations in the cyber context, compare the second and third row of Table 4.

[35] Extreme-value copulas allow to capture the dependence structure between certain rare events, for details see, e.g. Mai and Scherer (2017). The necessity of dealing adequately with *extreme* events in the cyber context has been emphasised by many authors, e.g. the comprehensive data-driven analysis of cyber losses by Eling and Wirfs (2019) advocated for distinguishing between "cyber risks of daily life" and "extreme cyber risks".

[36] Note that generally, an Archimedean copula is not parametrised by a parameter θ, but by the so-called (Archimedean) generator $\psi = \psi_\theta$, a non-increasing function $\psi : [0, \infty) \to [0, 1]$ with $\psi(0) = 1$ and $\lim_{x\to\infty} \psi(x) = 0$. The Gumbel copula is obtained by using the parametric family $\psi_\theta(x) = \exp\Big(-x^{\frac{1}{\theta}}\Big)$; for brevity, we use the notation C_θ instead of C_{ψ_θ}.

As remarked above, however, the main drawback of such a top-down modelling approach is that it is not based on an attempt to causally understand the dependence between cyber losses; instead, its merit is based on the analytical decomposition in Theorem 2 (see section A.2 in the electronic supplementary information) and its tractability in statistical inference. This is a somewhat questionable foundation in the cyber context due to scarcity, limited reliability, and suspected non-stationarity of available data, limiting the informativeness of models estimated on past data for the prediction of future losses. Therefore, we do not go into more detail on this example, but reiterate that in principle it provides the same flexibility regarding the effect of risk reduction services in insurance policies as the examples treated in detail above.

Conclusion

In recent years, with demand for cyber insurance increasing tremendously, cyber insurance markets around the world have been growing and the range of available cyber policies has been continuously expanding. As policies continue to mature, many prospective insurance buyers and external cyber experts agree that pure risk transfer cannot be an optimal cyber-risk management solution. Instead, companies – insured or not – have to make ongoing efforts to keep their cybersecurity measures up-to-date, given the evolving cyber-threat landscape. Therefore, there is mutual benefit (for all stakeholders) in the combination of risk transfer and risk reduction measures, leading to the (prospective) ubiquitous offering of pre-incident and post-incident services.

In this study, we have dealt with this combination of risk reduction and risk transfer in the cyber insurance context, and in particular addressed the question of how such risk reduction services should be optimally priced from an insurer's viewpoint. We have illustrated how common services within cyber insurance can be classified into the concepts of self-protection and self-insurance, and have argued how insurers should make use of their unique position regarding knowledge about the current cyber-loss landscape to offer additional pre-incident (warning) services to their policyholders.

We have shown that in the univariate case, i.e. when pricing a single contract alone, an insurer using a distortion risk measure with concave distortion (i.e. being risk-neutral or risk-averse) never has an economic incentive to subsidise pure self-protection services (i.e. only considering the effect on loss probability, factoring out potential cross-effect on loss sizes and the prospect of gaining additional information) and will, thus, always shift their full cost to the insurance buyer. Interestingly, this does not generally hold for the pricing of self-insurance services or when taking a multivariate (portfolio) viewpoint, in which case it can be optimal (and in some cases even mandatory to find an acceptable contract for both parties) to share the cost of risk reduction service between insurer and policyholder. We illustrate this finding using toy examples of two risks with dependence mechanisms representative for the cyber context and one exemplary extension to a larger multivariate setting.

From the insurance buyers' point of view, the study serves to illustrate how their initial risk (when approaching the insurer) and their choice of (distortion) risk measure as well as the existing portfolio of the insurer can influence the insurance price offered to them for different contracts (i.e. how much risk reduction is implicitly incentivised for them by the insurer's choice of price structure).

Some interesting aspects, however, remain for future research. We restricted the insurance buyer's options to full or no insurance (as is customary for primary insurance in the cyber context), but one could extend this to more general payout functions (e.g. proportional at any share $\alpha \in [0, 1]$ or excess-of-loss per risk at different priorities and limits).[37] Furthermore, we have mentioned that in the cyber context, part of the risk should be considered non-insurable (e.g. reputational risk), an aspect that could generalize the modelling of the insurance buyer's optimisation problem.

From the insurer's point of view, the pricing of self-protection and self-insurance services has been studied disjointly, whereas in practice, the combination of both types of services within a policy is customary. Furthermore, we have only illustrated the insurer's portfolio viewpoint in bivariate examples and an exchangeable extension. Fully exploring the question of optimal offering of cyber services using an insurer's more general multivariate viewpoint on a portfolio of dependent policyholders comprises many interesting questions for future work.

Furthermore, especially due to the potential for extreme cyber losses resulting from single large losses or accumulation risk from a large cyber event, many insurers work with reinsurance providers to limit their exposure and manage their portfolio risk. This opens the potential to analyse a suitable Stackelberg game between insurer and reinsurer(s) or even a set-up involving all three parties (insurance buyer(s), insurer, and reinsurer(s)). In this context, also interesting questions about optimal risk sharing arise.

Lastly, we have argued that the understanding of the dependence between cyber losses is crucial for insurers, as purely top-down dependence modelling approaches may not be suitable in the highly dynamic, non-stationary cyber domain. Therefore, more empirical research on the dependence structures underlying cyber risk, e.g. to more accurately determine underlying common factors leading to simultaneous exposure to a certain cyber event, is certainly necessary to better understand the evolving cyber-threat landscape. Lastly, it should be mentioned that many related questions from a not purely mathematical viewpoint arise. For example, economically and legally, it needs to be investigated how to ideally set up cyber insurance policies including services such that all parties (insurer, insureds, and IT security experts as service providers) draw synergies from the collaboration. From a technical viewpoint, one important issue is how to effectively quantify (and monitor) the IT security landscape of a potentially highly complex enterprise for actuarial applications. These issues emphasise the importance of interdisciplinary collaboration and research in the cyber

[37] An immediate generalization is a proportional insurance share $\alpha \in \{0\} \cup [\alpha_0, 1]$, which could illustrate not only the two cases *no insurance* and *full insurance*, but additionally the case where the insurance buyer purchases a minimum feasible share of risk transfer α_0 in order to benefit from the risk reduction services within insurance; in other words, insurers could sell policies that customers would not buy from a pure risk transfer viewpoint by including attractive services.

insurance domain in order to tackle this challenging risk. This article is complemented by an electronic supplement (Appendix) containing a seminal discussion of risk-assessment services, mathematical preliminaries, proofs, case studies and extended calculations.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1057/s41288-023-00289-7.

Acknowledgements The first author would like to thank the Cyber Insurance department of ERGO Group AG (Germany) for the opportunity to join their team to gain insight into the cyber insurance industry for three months while working on this research. We furthermore thank the two anonymous referees whose feedback greatly helped to improve the quality and presentation of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

Advisen. 2021. Advisen CyberGuide. https://cyberguide.advisenltd.com/.

Advisen and PartnerRe. 2017. Survey of cyber insurance market trends. https://www.partnerre.com/wp-content/uploads/2017/10/PartnerRe-2017-Survey-of-CyberInsurance-Market-Trends.pdf.

Advisen and PartnerRe. 2018. Survey of cyber insurance market trends. https://www.partnerre.com/wp-content/uploads/2018/10/2018-Survey-of-CyberInsurance-Market-Trends.pdf.

Allianz Global Corporate & Specialty SE. 2020. Allianz risk barometer—identifying the major business risks for 2020. https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/risk-barometer/Allianz-Risk-Barometer-2020-Risiken-Global.jpg.

Anchen, J., and D. Pain. 2017. Cyber: getting to grips with a complex risk. *Sigma* (1).

Artzner, P., F. Delbaen, J. Eber, and D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* 9 (3): 203–228.

Awiszus, K., Knispel, T., Penner, I. et al. 2023. Modeling and pricing cyber insurance. *European Actuarial Journal.*

Bai, Y., Z. Zhou, H. Xiao, R. Gao, and F. Zhong. 2022. A hybrid stochastic differential reinsurance and investment game with bounded memory. *European Journal of Operational Research* 296 (2): 717–737.

Balbás, A., B. Balbás, and A. Heras. 2011. Stable solutions for optimal reinsurance problems involving risk measures. *European Journal of Operational Research* 214 (3): 796–804.

Bellini, F., and C. Caperdoni. 2007. Coherent distortion risk measures and higher-order stochastic dominances. *North American Actuarial Journal* 11 (2): 35–42.

Bensalem, S., N. Hernández Santibáñez, and N. Kazi-Tani. 2020. Prevention efforts insurance, demand and price incentives under coherent risk measures. *Insurance: Mathematics and Economics* 93 (3): 369–386.

Bettels, S., S. Kim, and S. Weber. 2022. Multinomial backtesting of distortion risk measures. arXiv:2201.06319.

Bleichrodt, H., and L. Eeckhoudt. 2006. Willingness to pay for reductions in health risks when probabilities are distorted. *Health economics* 15 (2): 211–214.

Böhme, R. 2005. Cyber insurance revisited. In *Proceedings of the fourth workshop on the economics of information security (WEIS 2005)*. http://infosecon.net/workshop/pdf/15.pdf

Böhme, R., and G. Kataria. 2006. Models and measures for correlation in cyber insurance. In *WEIS*, vol 2.

Böhme, R., S. Laube, and M. Riek. 2008. A fundamental approach to cyber risk analysis. *Variance* 11(2).

Boonen, T., K. Cheung, and Y. Zhang. 2021. Bowley reinsurance with asymmetric information on the insurer's risk preferences. *Scandinavian Actuarial Journal* 2021 (7): 623–644.

Boonen, T., and M. Ghossoub. 2022. Bowley vs. pareto optima in reinsurance contracting. *European Journal of Operational Research*.

Boonen, T., and Y. Zhang. 2022. Bowley reinsurance with asymmetric information: a first-best solution. *Scandinavian Actuarial Journal* 2022 (6): 532–551.

Born, G. ProxyShell: Massive Angriffswelle auf ungepatchte Exchange-Server. https://www.heise.de/news/ProxyShell-Massive-Angriffswelle-auf-ungepatchte-Exchange-Server-6171597.html. Accessed 22 Aug 2021.

Bosch CyberCompare. https://cybercompare.com/. Accessed 20 Jan 2023.

Chase, J., D. Niyato, P. Wang, S. Chaisiri, and R. Ko. 2017. A scalable approach to joint cyber insurance and security-as-a-service provisioning in cloud computing. *IEEE Transactions on Dependable and Secure Computing* 16 (4): 565–579.

Chen, L., and Y. Shen. 2018. On a new paradigm of optimal reinsurance: a stochastic stackelberg differential game between an insurer and a reinsurer. *ASTIN Bulletin* 48 (02): 905–960.

Chen, L., Y. Shen, and J. Su. 2020. A continuous-time theory of reinsurance chains. *Insurance: Mathematics and Economics* 95 (3): 129–146.

Cheung, K., S. Yam, and Y. Zhang. 2019. Risk-adjusted Bowley reinsurance under distorted probabilities. *Insurance: Mathematics and Economics* 86: 64–72.

Christoffersen, P., and D. Pelletier. 2004. Backtesting value-at-risk: a duration-based approach. *Journal of Financial Econometrics* 2 (1): 84–108.

Courbage, C. 2001. Self-insurance, self-protection and market insurance within the dual theory of choice. *The Geneva Papers on Risk and Insurance Theory* 26 (1): 43–56.

Courbage, C., B. Rey, and N. Treich. 2013. Prevention and precaution. In *Handbook of insurance*. vol. 21, 185–204. New York: Springer.

Deelstra, G., J. Dhaene, and M. Vanmaele. 2011. An overview of comonotonicity and its applications in finance and insurance. In *Advanced mathematical methods for finance*, 155–179.

Denneberg, D. 2013. *Non-additive measure and integral*, vol. 27. Springer Science & Business Media

Dhaene, J., A. Kukush, D. Linders, and Q. Tang. 2012. Remarks on quantiles and distortion risk measures. *European Actuarial Journal* 2 (2): 319–328.

Dhaene, J., S. Vanduffel, M. Goovaerts, R. Kaas, Q. Tang, and D. Vyncke. 2006. Risk measures and comonotonicity: a review. *Stochastic Models* 22 (4): 573–606.

Dionne, G., and L. Eeckhoudt. 1985. Self-insurance, self-protection and increased risk aversion. *Economics Letters* 17 (1–2): 39–42.

Dou, W., W. Tang, X. Wu, L. Qi, X. Xu, X. Zhang, and C. Hu. 2020. An insurance theory based optimal cyber insurance contract against moral hazard. *Information Sciences* 527: 576–589.

Ehrlich, I., and G. Becker. 1972. Market insurance, self-insurance, and self-protection. *Journal of Political Economy* 80 (4): 623–648.

Eling, M., and K. Jung. 2018. Copula approaches for modeling cross-sectional dependence of data breach losses. *Insurance: Mathematics and Economics* 82: 167–180.

Eling, M., and J. Wirfs. 2019. What are the actual costs of cyber risk events? *European Journal of Operational Research* 272 (3): 1109–1119.

Embrechts, P., H. Liu, and R. Wang. 2018. Quantile-based risk sharing. *Operations Research* 66 (4): 936–949.

Emmer, S., D. Tasche, and M. Kratz. 2015. What is the best risk measure in practice? A comparison of standard measures. *Journal of Risk* 18 (2): 31–60.

ENISA. 2021. ENISA THREAT LANDSCAPE 2021: April 2020 to mid-July 2021. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2021.

Epstein, L. 1980. Decision making and the temporal resolution of uncertainty. *International Economic Review*, 269–283.

Etner, J., and M. Jeleva. 2013. Risk perception, prevention and diagnostic tests. *Health economics* 22 (2): 144–156.

Fahrenwaldt, M., S. Weber, and K. Weske. 2018. Pricing of cyber insurance contracts in a network model. *ASTIN Bulletin* 48 (3): 1175–1218.

Föllmer, H., and A. Schied. 2016. *Stochastic finance: An Introduction in Discrete Time*. Berlin: De Gruyter.

Fudenberg, D., and J. Tirole. 1991. *Game theory*. Cambridge: MIT Press.

Galchion, A. 2010. The V@R at risk. *International Journal of Theoretical and Applied Finance* 13 (4): 503–506.

Gordon, L.A., and M. Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* 5 (4): 438–457.

Gumbel, E.J. 1960. Bivariate exponential distributions. *Journal of the American Statistical Association* 55 (292): 698–707.

Herath, H., and T. Herath. Copula-based actuarial model for pricing cyber insurance policies. Insurance Markets and Companies, 2(1), 2011.

Hiebert, L. 1989. Optimal loss reduction and increases in risk aversion. *The Journal of Risk and Insurance* 56 (2): 300.

Holmstrom, B. 1979. Moral hazard and observability. *The Bell Journal of Economics* 10 (1): 74.

Jewitt, I. 1988. Justifying the first-order approach to principal-agent problems. *Econometrica* 56 (5): 1177.

Khalili, M., M. Liu, and S. Romanosky. 2019. Embracing and controlling risk dependency in cyber insurance policy underwriting. *Journal of Cybersecurity* 5 (1): 519.

Khalili, M., P. Naghizadeh, and M. Liu. 2018. Designing cyber insurance policies: the role of pre-screening and security interdependence. *IEEE Transactions on Information Forensics and Security* 13 (9): 2226–2239.

Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73 (6): 1849–1892.

Konrad, K., and S. Skaperdas. 1993. Self-insurance and self-protection: a nonexpected utility analysis. *The Geneva Papers on Risk and Insurance Theory* 18 (2): 131–146.

Kratz, M., Y. Lok, and A. McNeil. 2018. Multinomial var backtests: a simple implicit approach to backtesting expected shortfall. *Journal of Banking & Finance* 88: 393–407.

Kusuoka, S. 2001. On law invariant coherent risk measures. In *Advances in mathematical economics*, vol. 3, 83–95. Berlin: Springer

Lelarge, M., and J. Bolot. 2009. Economic incentives to increase security in the internet: the case for insurance. In *IEEE INFOCOM 2009—the 28th conference on computer communications*, 1494–1502.

Liu, S., and Q. Zhu. 2022. Mitigating moral hazard in cyber insurance using risk preference design. arXiv:2203.12001

Mai, J., and M. Scherer. 2017. *Simulating copulas: stochastic models, sampling algorithms, and applications*, vol. 6. Singapore: World Scientific Publishing.

Marotta, A., F. Martinelli, S. Nanni, A. Orlando, and A. Yautsiukhin. 2017. Cyber insurance survey. *Computer Science Review* 24 (3): 35–61.

Marshall, A., and I. Olkin. 1967. A multivariate exponential distribution. *Journal of the American Statistical Association* 62 (317): 30–44.

Mazzoccoli, A., and M. Naldi. 2020. Robustness of optimal investment decisions in mixed insurance/investment cyber risk management. *Risk Analysis* 40 (3): 550–564.

Mazzoccoli, A., and M. Naldi. 2021. Optimal investment in cybersecurity under cyber insurance for a multi-branch firm. *Risks* 9 (1): 24.

Müller, A., and D. Stoyan. 2002. *Comparison methods for stochastic models and risks*, vol. 389. New York: Wiley.

Munich, R. 2021. Cyber insurance: risks and trends. 2021. https://www.munichre.com/topics-online/en/digitalisation/cyber/cyberinsurance-risks-and-trends-2021.html.

Niculescu, C., and L. Persson. 2018. *Convex functions and their applications: a contemporary approach*. Berlin: Springer.

Osborne, M., and A. Rubinstein. 1994. *A course in game theory*. Cambridge: MIT Press.

Pal, R., and L.Golubchik. 2010. Analyzing self-defense investments in internet security under cyber-insurance coverage. In *2010 IEEE 30th international conference on distributed computing systems*, 339–347. IEEE

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2014. Will cyber insurance improve network security? A market analysis. In *IEEE INFOCOM 2014-IEEE conference on computer communications*, 235–243. IEEE

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2017. Security pricing as enabler of cyber insurance a first look at differentiated pricing markets. *IEEE Transactions on Dependable and Secure Computing* 16 (2): 358–372.

Parkin, M., M. Powell, and K. Matthews. 2002. *Economics*. Harlow: Addison-Wesley.

Peng, C., M. Xu, S. Xu, and T. Hu. 2018. Modeling multivariate cybersecurity risks. *Journal of Applied Statistics* 45 (15): 2718–2740.

Rogerson, W. 1985. The first-order approach to principal-agent problems. *Econometrica* 53 (6): 1357.

Schwartz, G., and S. Sastry. 2014. Cyber insurance framework for large scale interdependent networks. In *Proceedings of the 3rd international conference on High confidence networked systems*, 145–154. ACM.

Schwartz, G., N. Shetty, and J. Walrand. 2013. Why cyber insurance contracts fail to reflect cyber-risks. In *2013 51st annual Allerton conference on communication, control, and computing*, 781–787. IEEE.

Shetty, S., M. McShane, L. Zhang, J. Kesan, C. Kamhoua, K. Kwiat, and L. Njilla. 2018. Reducing informational disadvantages to improve cyber risk management. *The Geneva Papers on Risk and Insurance—Issues and Practice* 43 (2): 224–238.

Shetty, N., G. Schwartz, M. Felegyhazi, and J. Walrand. 2010. Competitive cyberinsurance and internet security. In *Economics of Information Security and Privacy*, vol 5, 229–247. Berlin: Springer.

Shetty, N., G. Schwartz, and J. Walrand. 2010. Can competitive insurers improve network security? In *Trust and Trustworthy Computing*. Lecture Notes in Computer Science, vol. 610, 1308–322. Heidelberg: Springer.

Skeoch, H. 2022. Expanding the gordon-loeb model to cyber insurance. *Computers & Security* 112: 102533.

Sklar, A. 1959. Fonctions de repartition à n dimensions et leurs marges. *Publications de l'Institut de statistique de l'Université de Paris* 8: 229–231.

tenable. 2021. *tenable's 2021 threat landscape retrospective*. tenable Research

Vaserstein, L. 1969. Markov processes over denumerable products of spaces, describing large systems of automata. *Problemy Peredachi Informatsii* 5 (3): 64–72.

Wang, S. 1995. Insurance pricing and increased limits ratemaking by proportional hazards transforms. *Insurance: Mathematics and Economics* 17 (1): 43–54.

Wang, S. 1996. Premium calculation by transforming the layer premium density. *ASTIN Bulletin: The Journal of the IAA* 26 (1): 71–92.

Wang, S. 1998. Implementation of proportional hazards transforms in ratemaking. *Proceedings of the Casualty Actuarial Society* 85 (1–2): 940–979.

Wang, S. 2000. A class of distortion operators for pricing financial and insurance risks. *The Journal of Risk and Insurance* 67 (1): 15.

Wang, R. 2016. Regulatory arbitrage of risk measures. *Quantitative Finance* 16 (3): 337–347.

Wang, S., V. Young, and H. Panjer. 1997. Axiomatic characterisation of insurance prices. *Insurance: Mathematics and economics* 21 (2): 173–183.

Weber, S. 2018. Solvency II, or how to sweep the downside risk under the carpet. *Insurance: Mathematics and economics* 82: 191–200.

Wirch, J., and M. Hardy. 1999. A synthesis of risk measures for capital adequacy. *Insurance: Mathematics and Economics* 25 (3): 337–347.

Wirch, J., and M. Hardy. 2000. Ordering of risk measures for capital adequacy. Institute of Insurance and Pension Research, University of Waterloo, Research Report 00–03.

Xiang, Q., A. Neufeld, G. Peters, I. Nevat, and A. Datta. 2021. A bonus-malus framework for cyber risk insurance and optimal cybersecurity provisioning. arXiv:2102.05568.

Xu, M., G. Da, and S. Xu. 2015. Cyber epidemic models with dependences. *Internet Mathematics* 11 (1): 62–92.

Xu, M., and L. Hua. 2019. Cybersecurity insurance: modeling and pricing. *North American Actuarial Journal* 23 (2): 220–249.

Yang, Z., and J. Lui. 2014. Security adoption and influence of cyber insurance markets in heterogeneous networks. *Performance Evaluation* 74: 1–17.

Yin, T., A. Sarabi, and M. Liu. 2021. Deterrence, backup, or insurance: a game-theoretic analysis of ransomware. In *The Annual Workshop on the Economics of Information Security* (*WEIS*).

Young, D., J. Lopez, M. Rice, B. Ramsey, and R. McTasney. 2016. A framework for incorporating insurance in critical infrastructure cyber risk strategies. *International Journal of Critical Infrastructure Protection* 14: 43–57.

Zeller, G., and M. Scherer. 2022. A comprehensive model for cyber risk based on marked point processes and its application to insurance. *European Actuarial Journal* 12 (1): 33-85.

Zhang, R., and Q. Zhu. 2021. Optimal cyber insurance contract design for dynamic risk management and mitigation. *IEEE Transactions on Computational Social Systems*.

Zhang, R., Q. Zhu, and Y. Hayel. 2017. A bi-level game approach to attack-aware cyber insurance of computer networks. *IEEE Journal on Selected Areas in Communications* 35 (3): 779–794.

Ziggel, D., T. Berens, G. Weiß, and D. Wied. 2014. A new set of improved value-at-risk backtests. *Journal of Banking & Finance* 48: 29–41.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

About the authors

Gabriela Zeller is a doctoral candidate at ERGO Center of Excellence in Insurance at the Chair of Mathematical Finance at the Technical University of Munich (TUM). Before starting her PhD, she graduated from the elite graduate program Finance and Information Management at TUM and the Universities of Augsburg and Bayreuth with a thesis on Hawkes processes in insurance. She has completed study and research stays in Mexico, Sweden and Canada.

Matthias Scherer is a professor for Risk and Insurance at the Technical University of Munich. His research interests comprise the pricing and risk management of financial derivatives, probability theory, statistics and efficient numerical tools. He is particularly interested in copula models and multivariate financial problems.

A Electronic Supplementary Information

This is the electronic supplement to the article: Zeller, G. and M. Scherer. 2023. *Risk mitigation services in cyber insurance: optimal contract design and price structure.*
It contains a seminal discussion of risk-assessment services, mathematical preliminaries, proofs, case studies, and extended calculations.

A.1 A Note on Risk-Assessment Services

An important, yet challenging, aspect of cyber insurance is risk assessment. In established insurance lines, insurers can rely on plenty of historical claims and in-house expertise to classify prospective policyholders into risk categories (e.g. using standardized models for mass segments such as motor insurance, individual expert judgement for large risks in industrial lines, or a mixture of both). In contrast, due to the novelty, non-stationarity, and complexity of cyber risk, the choice of relevant risk factors and the process of their assessment is a challenging task. For example, it seems intuitive that a policyholder's IT infrastructure and existing cyber security provisions are an important factor for his susceptibility to cyber risk. However, how to include extensive qualitative information about a system's structure and potential vulnerabilities into an actuarial model is a complex open issue in itself. Nevertheless, insurers often cooperate with specialized IT security experts to conduct extensive IT audits of a prospective policyholder before pricing insurance coverage (and choosing to make an offer at all). While the exact nature of these collaborations is of course opaque from the outside, it is likely that insurers currently carry the full costs of these audits and subsume them under operational or acquisition costs of their cyber insurance business. Thus, for such services, the following questions naturally arise: *Which amount of effort should an insurer optimally invest in risk assessment? Under which circumstances is it (feasible and) optimal for an insurer to shift part of the cost of risk assessment to the insurance buyer?* To elaborate on these questions, three relevant points should be mentioned:

- An extensive risk assessment benefits both parties: While the insurer increases her knowledge about the to-be-insured risk (allowing for more accurate pricing and risk management), the prospective insurance buyer also benefits from gaining expert knowledge about his IT security situation, often even including an ordered list of priorities to be addressed to ensure efficient spending of limited IT (security) budget. This is particularly beneficial for small to medium-size enterprises without sophisticated internal IT divisions, for whom the risk assessment as part of the insurance take-up process may be the first comprehensive analysis of the cyber-security level of their organization. While not every inquiry about insurance prices leads to the closure of a cyber-insurance contract, the process may serve as a wake-up call for the acquisition of (additional) risk reduction measures within or outside of an insurance policy.
- It is clear that due to the dynamics of the cyber threat landscape and the ongoing evolution of IT ecosystems, the risk assessment process should not be conducted only once, but periodically during the life (and at the renewal) of a cyber-insurance policy. From this perspective, one may argue that this is a service offering complementary to the ones discussed in the main body of this study. While *self-protection* and *self-insurance* services

aim to reduce the loss probability and loss severity respectively, i.e. re-shape the original loss distribution (previously denoted F_0) favourably, *risk assessment* services acknowledge that neither party has accurate knowledge of F_0 and therefore aim at better understanding it.

- In our view, the issue of *incomplete information* in the cyber context is not exactly akin to *asymmetric information* or *moral hazard.* By that we mean that in reality, it is not usually the case that an insurance buyer has better knowledge (in a meaningful, quantifiable way) about their own IT systems than the insurer and seeks to obscure this knowledge from the insurer, such that she has to use risk assessment to extract this hidden information. Rather, both parties jointly seek the expertise of an IT security provider to better understand the underlying risk and thus through the risk assessment process jointly gain improved (shared) knowledge. Apart from the insurance buyer's own lack of information, it has to be considered that as cyber risks are partly non-insurable (in particular regarding reputational risks) and cyber policies (as other lines) come with exclusions in case relevant information is wrongly stated or willfully omitted at policy closure, the insurance buyer has no incentive to take malicious advantage of asymmetric information.

There are, so far, very few mathematical papers concerned with cyber risk assessment services.[39] Below, we propose a general mathematical framework to describe the setting where none of the parties has full information, but both are able to learn more about the true underlying distribution through a (costly) risk audit. The further study of this setting is outside the scope of this manuscript but remains an interesting problem for future research.

0. As before, denote by F_0 the *true* (unknown) distribution of the risk X the buyer faces.
1. As the buyer approaches the insurer, both parties have (distinct) subjective views about the distribution of the risk X:

   $(Buyer : pre\ audit)$: $F_{0,B}$ depending on his own (incomplete) knowledge about his IT systems.

   $(Insurer : pre\ audit)$: $F_{0,I}$ depending on her standard models for cyber and publicly available information about the buyer.

   Note that while insurers routinely think in terms of risk and probability distributions, already assuming that a regular insurance buyer would have a certain distribution in mind is one step ahead of reality. In practice, most companies may consider scenarios and a setting of uncertainty rather than an already formalized context of risk.
2. If an IT audit is conducted, this causes some effort $a > 0$ (where different options w.r.t. extensiveness are usually available) with associated cost $c(a)$ (with the typical assumptions

[39] An exception is [Khalili et al., 2018], who study the role of pre-screening (i.e. risk assessment) in cyber insurance (again in a setting of *asymmetric* information, i.e. assuming the insurance buyer has perfect information). They study the influence of risk aversion (in an expected utility setting) and interdependence of insurance buyers under very specific assumptions on the loss distribution (normal distribution with parameters depending on the insurance buyer's effort), the effect of effort on the loss distribution, the pre-screening process (the outcome being the true effort plus standard Gaussian noise), and the interdependence between agents.

of increasingness and convexity). At the conclusion of the audit, an IT security expert shares their subjective opinion with all other parties:

$$(\textit{IT Expert}:\ \textit{audit})\ :\ F^a_{0,A}\ \text{ depending on the extent of the audit } a > 0.$$

The more effort is invested in the audit, the closer the expert assessment approaches the *true* distribution, i.e.

$$0 < a_1 \le a_2 < \infty :\ \ \mathcal{D}\big(F^{a_1}_{0,A}, F_0\big) \ge \mathcal{D}\big(F^{a_2}_{0,A}, F_0\big) \ge 0, \tag{12}$$

where $\mathcal{D}$ denotes some measure of distance between probability distributions (e.g. the Wasserstein distance between the associated probability measures, see e.g. [Vaserstein, 1969]).

3. Both buyer and insurer update their original views by incorporating the result of the IT audit, i.e.

$$\begin{aligned}(\textit{Buyer}:\textit{post audit})\ \ \hat{F}_{0,B}(a) &:= h_B(F_{0,B}, F^a_{0,A}) \overset{e.g.}{=} w_B F_{0,B} + (1-w_B)F^a_{0,A}, \quad w_B \in [0,1],\\ (\textit{Insurer}:\textit{post audit})\ \ \hat{F}_{0,I}(a) &:= h_I(F_{0,I}, F^a_{0,A}) \overset{e.g.}{=} w_I F_{0,I} + (1-w_I)F^a_{0,A}, \quad w_I \in [0,1].\end{aligned}$$

   As an example, we have assumed that each party uses a convex combination (mixture) of their original view and the new information, with potentially different weights. In general, $h_B(\cdot)$ and $h_I(\cdot)$ could denote any functions mapping two distribution functions (*a priori* and *IT audit*) to a new distribution function (*a posteriori*).

4. As an initial step, one may formulate the insurer's loss function depending on the risk assessment effort as

$$\begin{aligned}(\textit{pre audit})\ \ &\Big[\rho_{I,F_{0,I}}(X) - (1+\theta)\mathbb{E}_{F_{0,I}}[X]\Big]\mathbb{1}_{\{(1+\theta)\mathbb{E}_{F_{0,I}}[X]\le\rho_{B,F_{0,B}}(X)\}},\\ (\textit{post audit})\ \ &\Big[\rho_{I,\hat{F}_{0,I}(a)}(X) - (1+\theta)\mathbb{E}_{\hat{F}_{0,I}(a)}[X]\Big]\mathbb{1}_{\{(1+\theta)\mathbb{E}_{\hat{F}_{0,I}(a)}[X]\le\rho_{B,\hat{F}_{0,B}(a)}(X)\}} + c(a),\end{aligned}$$

   where $\rho_{I,\cdot}, \rho_{B,\cdot}$ denote the risk measures of the risk X used by the insurer and buyer, respectively, given their subjective views about the distribution of X (*a priori* without audit and *a posteriori* with audit).

However, the above loss function does not yet encode the insurer's preference for a better estimation of the true distribution, i.e. the gain resulting from (12). Crucially, note that in this setting, we are in the framework of *precaution* and *uncertainty* as opposed to *prevention* and *risk* in the main body of the paper. An excellent summary of both concepts and the corresponding literature can be found in [Courbage et al., 2013]. As outlined therein, there are several approaches to formalizing the *effect of more information*; a classical result (without intertemporal dependence) in an expected utility setting is due to the seminal work of [Epstein, 1980] whose results have been used as the foundation for many following studies. Another (contested) approach to *precaution* is the idea of ambiguity (aversion), see [Klibanoff et al., 2005]. The choice and application of such an approach to the above problem is an interesting next step in a consecutive research project.

A.2 Mathematical Prerequisites

Definition 1 (Law-invariant, coherent risk measure (see [Artzner et al., 1999, Kusuoka, 2001]))**.** *A risk measure $\rho : L^1(\Omega, \mathcal{F}, \mathbb{P}) \to \mathbb{R}$ is a coherent risk measure if it has the following properties:*

1. *Monotonicity: $X \leq Y \implies \rho(X) \leq \rho(Y)$;*
2. *Cash-additivity / translation invariance: $\forall\, m \in \mathbb{R} :\ \rho(X - m) = \rho(X) - m$;*
3. *Convexity: $\forall\, X, Y \in L^1, \forall\, \lambda \in [0,1] :\ \rho(\lambda X + (1-\lambda)Y) \leq \lambda\rho(X) + (1-\lambda)\rho(Y)$;*
4. *Positive homogeneity: $\forall\, \lambda \in [0,\infty) :\ \rho(\lambda X) = \lambda\rho(X)$.*

Note that under property 4., convexity is equivalent to sub-additivity. A coherent risk measure is called law-invariant if additionally:
5. Law-invariance: $X \overset{(d)}{=} Y \implies \rho(X) = \rho(Y)$, where $\overset{(d)}{=}$ denotes equality in distribution.

Definition 2 (Copula, see, e.g., [Mai and Scherer, 2017])**.** *A function $C : [0,1]^d \to [0,1]$ is called copula, if there is a random vector $(U_1, \ldots, U_d)$ such that $U_j \sim \mathcal{U}[0,1],\ j \in \{1, \ldots, d\}$, and C is the joint c.d.f. of $(U_1, \ldots, U_d)$, i.e.*

$$C(u_1, \ldots, u_d) = \mathbb{P}(U_1 \leq u_1, \ldots, U_d \leq u_d),\ \ u_1, \ldots, u_d \in [0,1].$$

Theorem 2 (Sklar's Theorem, see [Sklar, 1959])**.** *A function $F : \mathbb{R}^d \to [0,1]$ is the c.d.f. of a random vector $(X_1, \ldots, X_d)$ iff there exist a copula $C : [0,1]^d \to [0,1]$ and univariate c.d.f.s $F_1, \ldots, F_d : \mathbb{R} \to [0,1]$ s.t.*

$$C\big(F_1(x_1), \ldots, F_d(x_d)\big) = F(x_1, \ldots, x_d),\ \ \forall x_1, \ldots, x_d \in \mathbb{R}.$$

A.3 Proofs and Derivations of Sections 2 and 3

The following lemma details the relationship between the parameter s and the distortion risk measure $\rho_s(X) := \int_0^\infty \psi(\overline{F}_{X,s}(x))\mathrm{d}x$ under the assumptions of Section 2.

Lemma 1 ([Bensalem et al., 2020], Lemma 2.1)**.** *Let ρ be a coherent, law-invariant risk measure and assume X to be distributed according to a family of distributions F_s such that* (A2) *holds. Then, the map $s \mapsto \rho_s(X)$ is convex, continuous, non-increasing and $\rho_s(X) \geq \mathbb{E}_s[X] > 0$.*

Proof. *See proof of Lemma 2.1 in [Bensalem et al., 2020].*

Choosing $\psi(u) = u$, the statements of Lemma 1 carry over to the special case $s \mapsto \mathbb{E}_s[X]$. For notational convenience, we formalize the following:

Assumption 3 (Monotonicity of $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$)**.** *Assume that $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$ is monotone non-decreasing, i.e.*

$$\forall\, 0 \leq s_1 \leq s_2 < \infty :\ \frac{\rho_{s_2}(X)}{\mathbb{E}_{s_2}[X]} \geq \frac{\rho_{s_1}(X)}{\mathbb{E}_{s_1}[X]} \iff 0 \leq \frac{\rho_{s_1}[X] - \rho_{s_2}[X]}{\rho_{s_1}[X]} \leq \frac{\mathbb{E}_{s_1}[X] - \mathbb{E}_{s_2}[X]}{\mathbb{E}_{s_1}[X]}. \tag{A3}$$

Lemma 2 (A sufficient condition for Assumption 3 - [Bensalem et al., 2020], Lemma 3.2). *Let ρ be a distortion risk measure with concave distortion function ψ and let the distribution of X be as in* (4)*. Then, the map $s \mapsto \frac{\rho_s(X)}{\mathbb{E}_s[X]}$ is non-decreasing.*

An alternative proof to the one given in [Bensalem et al., 2020] is given below.

Proof (Lemma 2). *To show* (A3)*, rearranging the equation and using that by Equation* (4) *the survival functions $\overline{F}_{X,s}$ are given by*

$$\overline{F}_{X,s}(x) = \begin{cases} 1 & \text{if } x < 0, \\ p(s)\overline{F}_Y(x) & \text{if } x \geq 0, \end{cases}$$

yield the need to show for $s_2 \geq s_1$

$$\rho_{s_2}(X) \geq \rho_{s_1}(X)\frac{\mathbb{E}_{s_2}[X]}{\mathbb{E}_{s_1}[X]} = \rho_{s_1}(X)\frac{p(s_2)\mathbb{E}[Y]}{p(s_1)\mathbb{E}[Y]} = \rho_{s_1}(X)\frac{p(s_2)}{p(s_1)}.$$

Note that for the concave function ψ with $\psi(0) = 0$ it holds

$$\forall t \in [0,1],\ x \in \mathbb{R}:\ \psi(tx) \geq t\psi(x), \tag{13}$$

implying

$$\begin{aligned} \rho_{s_1}(X)\frac{p(s_2)}{p(s_1)} &= \int_0^\infty \underbrace{\frac{p(s_2)}{p(s_1)}}_{\in[0,1]} \psi\Big(p(s_1)\overline{F}_Y(x)\Big)\mathrm{d}x \overset{(13)}{\leq} \int_0^\infty \psi\Big(\frac{p(s_2)}{p(s_1)}p(s_1)\overline{F}_Y(x)\Big)\mathrm{d}x \\ &= \int_0^\infty \psi\Big(p(s_2)\overline{F}_Y(x)\Big)\mathrm{d}x = \rho_{s_2}(X). \end{aligned}$$

Lemma 3 ([Bensalem et al., 2020], Lemma 3.3). *Assume that ρ is a distortion risk measure with a distortion function ψ such that $s \mapsto \psi(p(s))$ is convex. Then $s \mapsto \rho_s(X)$ and $s \mapsto \mathbb{E}_s[X]$ are convex.*

Proof (Lemma 3). *See proof of Lemma 3.3 in [Bensalem et al., 2020].*

Proof (Corollary 1). *The proof follows the derivation in [Bensalem et al., 2020], p.375 - where $s_B(\theta,\beta)$ for fixed $\beta \in [\underline{\beta}, 1]$ corresponds to e_θ therein - with the following amendments:*
1. The introduction of the constant θ_0 allows the inclusion of the values $s = 0$ and $s = s_B(\theta,\beta)$ in either of the sets $\mathcal{I}$ or $\mathcal{N}$, and the statement of increasingness (instead of non-decreasingness) in the case $\theta > \theta_0$.
2. To derive the definition of $s_B(\theta,\beta)$ and show its increasingness in θ, we use the map $G^\beta(s)$ with given assumptions on $s \mapsto c(s)$ (increasing and continuous with $c(0) = 0$ and $\lim\limits_{s\to\infty} c(s) = \infty$).
3. We need to show increasingness of $\beta \mapsto s_B(\theta,\beta)$:
Note that for any $s > 0$, $\beta \mapsto G^\beta(s)$ is decreasing: Let $\underline{\beta} \leq \beta_1 < \beta_2 \leq 1$, then

$$G^{\beta_1}(s) = \frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]} + \underbrace{(\beta_o - \beta_1)}_{>0}\underbrace{\frac{c(s)}{\mathbb{E}_s[X]}}_{>0} > \frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]} + (\beta_o - \beta_2)\frac{c(s)}{\mathbb{E}_s[X]} = G^{\beta_2}(s).$$

Therefore, $s_B(\theta,\beta)$ is increasing in β:

$$s_B(\theta,\beta_1) = \min\{s > 0 : G^{\beta_1}(s) \geq 1 + \theta\} < \min\{s > 0 : G^{\beta_2}(s) \geq 1 + \theta\} = s_B(\theta,\beta_2).$$

Proof (Corollary 2)**.**
Part 1.:
First of all, assume that $\theta > \theta_0$, else $\mathcal{N}$ is empty and a local minimizer on this set does not exist. Following the proof of Proposition 3.4 in [Bensalem et al., 2020], we define for fixed β the (finite) constant

$$\theta_N(\beta) := \inf\{\theta \geq 0 : s_B(\theta,\beta) > s_N\}$$

*to derive the statement, where we note that in the case $\theta < \theta_N(\beta)$, the argument "the concave map $L_{\mathcal{N}}$ is non-decreasing on $\mathcal{N}$" has to be replaced by "the **convex** map $L_{1,\mathcal{N}}$ is non-**in**creasing on $\mathcal{N}$".*
Furthermore, we note that as $\beta \mapsto s_B(\theta,\beta)$ is increasing by Lemma 1 and s_N does not depend on β, by definition $\beta \mapsto \theta_N(\beta)$ is decreasing.
Part 2.:
The definition of the constant $\theta_I(\beta)$ for any fixed $\beta \in [\underline{\beta}, 1]$ and the increasingness of $\theta \mapsto s_I(\theta,\beta)$ follow the first part of the proof of Proposition 3.2 in [Bensalem et al., 2020], whereby the loss function $L^{\theta,\beta}_{1,\mathcal{I}}$ is considered. We additionally have to show decreasingness of $\beta \mapsto s_I(\theta,\beta)$ in case $\theta > \theta_I(\beta)$. Here, the global minimizer $s_I(\theta,\beta)$ is an interior point on $(0,\infty)$ characterized by

$$[L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=s_I(\theta,\beta)} = 0 \iff \beta = \underbrace{\frac{1}{c'(s)|_{s=s_I(\theta,\beta)}}}_{>0,\ \textit{decreasing in } s} \underbrace{\left(-(1+\theta)\mathbb{E}'_s[X]|_{s=s_I(\theta,\beta)}\right)}_{>0,\ \textit{non-increasing in } s}. \tag{14}$$

As the left-hand side is increasing in β, so must be the right-hand side. As $s \mapsto \frac{1}{c'(s)}(-(1+\theta)\mathbb{E}'_s[X])$ is non-increasing as a product of two positive, (at least) non-increasing functions, the inner function $\beta \mapsto s_I(\theta,\beta)$ must be decreasing.

To prove Corollary 3, we introduce the following lemma, where for the rest of this section, we suppress the dependencies $s_I(\theta,\beta)$ and $s_B(\theta,\beta)$ for brevity.

Lemma 4 (Adapted from [Bensalem et al., 2020])**.** *If*

$$(N1)\ \ s_N < s_I < s_B \qquad or \qquad (N2)\ \ s_I \leq s_N < s_B,$$

the global minimizer of $L_1(\alpha^*(s), s)$ is $(\alpha^*, s^*) = (0, s_N)$.
If

$$(I1)\ \ s_B \leq s_I < s_N \qquad or \qquad (I2)\ \ s_B \leq s_N \leq s_I,$$

the global minimizer of $L_1(\alpha^*(s), s)$ is $(\alpha^*, s^*) = (1, s_I)$.
The case $s_I < s_B \leq s_N$ is not possible, i.e. the global minimizers of $L^{\theta,\beta}_{1,\mathcal{I}}(s)$ resp. $L_{1,\mathcal{N}}(s)$ cannot be simultaneously outside of $\mathcal{I}$ resp. $\mathcal{N}$.
If

$$(T)\ s_N < s_B \leq s_I,$$

for any $\beta \in [\underline{\beta}, 1]$, there exists a constant $\theta_R(\beta) \geq 0$ such that

(i) If $\theta \leq \theta_R(\beta)$, the global minimizer of $L_1(\alpha^*(s), s)$ is $(\alpha^*, s^*) = (1, s_I)$.

(ii) If $\theta > \theta_R(\beta)$, the global minimizer of $L_1(\alpha^*(s), s)$ is $(\alpha^*, s^*) = (0, s_N)$.

Furthermore, it holds $\theta_R(\beta) \geq \theta_N(\beta)$.

Proof (Lemma 4)**.** *See proof of Proposition 3.6 (for cases* $(N1)$ *and* $(N2)$*), Proposition 3.5/Corollary 3.2 (for cases* $(I1)$ *and* $(I2)$ *and the argument directly after), and Theorem 3.2 (for case* (T)*) in [Bensalem et al., 2020].*

Proof (Corollary 3)**.** *The statement follows by combining the results of Lemma 4 by considering the progression through the cases for* $\theta \in [0, \infty)$*.*

- $\theta = 0$ *implies* $s_B = 0$, $\mathcal{I} = [0, \infty)$*, thus at* $\theta = 0$*, either* $(I1)$ *or* $(I2)$ *holds.*
- *If at* $\theta = 0$, $(I1)$ *holds, a transition to* $(I2)$ *must occur as* θ *increases (as so do* s_I *and* s_B*), as the only alternative* $s_I < s_B \leq s_N$ *is impossible (see Lemma 4). As* s_I *is increasing in* θ*, transitioning back is not possible.*
- $(I1)$ *or* $(I2)$ *occur exactly while* $s_B \leq s_N$ *(meaning* $s_N \notin \mathcal{N}$*), i.e. for* $0 \leq \theta \leq \theta_N(\beta)$ *(see Corollary 2). The optimal solution is* $(\alpha^*, s^*) = (1, s_I)$*.*
- *After crossing* $\theta = \theta_N(\beta)$*, one transitions from* $s_B \leq s_N$ *in* $(I2)$ *to one of the cases with* $s_N < s_B$*, i.e.* (T) *or* $(N1)$*. Note that* $(N1)$ *cannot occur while* $\theta \leq \theta_M(\beta)$*, as by definition of* $\theta_M(\beta)$ *and global optimality of* s_N *for* $L_{1,\mathcal{N}}$*, for any* $\theta \leq \theta_M(\beta)$ *it holds that* $s_I \in \mathcal{I} = [s_B, \infty)$ *as*
$$L_{1,\mathcal{I}}^{\theta,\beta}(s_I) \leq L_{1,\mathcal{N}}(s_N) < L_{1,\mathcal{N}}(s_I).$$
- *This implies that if the transition at* $\theta_N(\beta)$ *is to* $(N1)$*, then* $\theta_M(\beta) = \theta_N(\beta)$ *and if the transition is to* (T)*, then* $\theta_M(\beta) > \theta_N(\beta)$*. In any case, for* $\theta \leq \theta_M(\beta)$*, implication* (i) *of Lemma 4 holds, i.e.* $(\alpha^*, s^*) = (1, s_I)$*. For* $\theta > \theta_M(\beta)$*, either implication* (ii) *of Lemma 4 or Lemma* $(N1)$ *holds; in either case,* $(\alpha^*, s^*) = (0, s_N)$*.*

The assertion that $\theta_R(\beta) \geq \theta_N(\beta)$ *follows from Lemma 4.*
It remains to show that $\beta \mapsto \theta_R(\beta)$ *is non-increasing. Note that for any* $\theta \geq 0$*, the map* $\beta \mapsto L_{1,\mathcal{I}}^{\theta,\beta}(s_I)$ *is non-decreasing:*
Let $\underline{\beta} \leq \beta_1 < \beta_2 \leq 1$*, then*

$$\begin{aligned}
L_{1,\mathcal{I}}^{\theta,\beta_1}\big(s_I(\theta,\beta_1)\big) &= (1+\theta)\mathbb{E}_{s_I(\theta,\beta_1)}[X] + \beta_1 c\big(s_I(\theta,\beta_1)\big) \\
&\leq (1+\theta)\mathbb{E}_{s_I(\theta,\beta_2)}[X] + \beta_1 c\big(s_I(\theta,\beta_2)\big) \\
&\leq (1+\theta)\mathbb{E}_{s_I(\theta,\beta_2)}[X] + \beta_2 c\big(s_I(\theta,\beta_2)\big) = L_{1,\mathcal{I}}^{\theta,\beta_2}\big(s_I(\theta,\beta_2)\big),
\end{aligned}$$

where the first inequality stems from the global optimality of $s_I(\theta,\beta_1)$ *for* $L_{1,\mathcal{I}}^{\theta,\beta_1}$ *and both inequalities are strict unless* $s_I(\theta,\beta_1) = s_I(\theta,\beta_2) = 0$*.*
This implies that $\beta \mapsto L_{1,\mathcal{I}}^{\theta,\beta}(s_I)$ *is constant for* $\theta \leq \theta_I(\beta)$ *and increasing for* $\theta > \theta_I(\beta)$*, thus non-decreasing. It follows from the definition of* $\theta_R(\beta)$ *that*

$$\begin{aligned}
\theta_R(\beta_1) &= \sup\{\theta \geq 0 : L_{1,\mathcal{I}}^{\theta,\beta_1}(s_I(\theta,\beta_1)) \leq L_{1,\mathcal{N}}(s_N)\} \\
&\geq \sup\{\theta \geq 0 : L_{1,\mathcal{I}}^{\theta,\beta_2}(s_I(\theta,\beta_2)) \leq L_{1,\mathcal{N}}(s_N)\} = \theta_R(\beta_2).
\end{aligned}$$

Proof (Corollary 4)**.** *See the proof of Corollary 3.3 in [Bensalem et al., 2020].*

A.4 Proofs of Section 4

Note that as above the derivative w.r.t. s will be denoted by the prime $()'$ and partial derivatives w.r.t. θ and β will be denoted explicitly as $\frac{\partial}{\partial\theta}$ and $\frac{\partial}{\partial\beta}$; occasionally, the dependency $s_I(\theta,\beta)$ is omitted for brevity. Furthermore, we assume all derivatives calculated in the following to exist.

Proof (Proposition 1). *The partial derivative of the insurer's loss w.r.t.* β *is given by*

$$\frac{\partial}{\partial\beta}L_0(\theta,\beta) = \underbrace{\frac{\partial}{\partial\beta}s_I(\theta,\beta)}_{\leq 0 \text{ by Cor. 2}}\Big[\underbrace{\rho'_{0,s}(X)|_{s=s_I} - (1+\theta)\mathbb{E}'_s[X]|_{s=s_I}}_{(*)} + \underbrace{(1-\beta)c'(s)|_{s=s_I}}_{\geq 0 \text{ by assumption}}\Big]\underbrace{-c(s_I(\theta,\beta))}_{\leq 0}$$

Showing $(*) > 0$ *directly implies* $\frac{\partial}{\partial\beta}L_0(\theta,\beta) \leq 0$ *such that the claim follows.*
Recall that a necessary condition for the insurer to be willing to offer a contract is that it entails a negative loss, i.e.

$$L_0(\theta,\beta) = \rho_{0,s_I(\theta,\beta)}(X) - (1+\theta)\mathbb{E}_{s_I(\theta,\beta)}[X] + (1-\beta)c(s_I(\theta,\beta)) < 0.$$

As the last term is non-negative, this implies the necessity of

$$\rho_{0,s_I(\theta,\beta)}(X) < (1+\theta)\mathbb{E}_{s_I(\theta,\beta)}[X], \tag{15}$$

i.e. that the insurer's measure of the risk taken over must be at least compensated by the premium received (and unless $\beta = 1$*, the difference must furthermore compensate the additional cost taken over by the provision of services). Recall that for a risk-averse insurer, Lemma 2 states that for any* $s \geq 0$

$$0 \geq \mathbb{E}_s[X]\rho'_s(X) \geq \rho_s(X)\mathbb{E}'_s[X] \iff 0 \leq \mathbb{E}_s[X](-\rho'_s(X)) \leq \rho_s(X)(-\mathbb{E}'_s[X]),$$

which, together with Equation (15)*, implies that in particular also at* $s = s_I$*,*

$$0 \leq \mathbb{E}_{s_I}[X](-\rho'_{0,s_I}(X)) \leq \rho_{0,s_I}(X)(-\mathbb{E}'_{s_I}[X]) \overset{(15)}{<} (1+\theta)\mathbb{E}_{s_I}[X](-\mathbb{E}'_{s_I}[X]).$$

Dividing by $\mathbb{E}_{s_I}[X] > 0$ *and rearranging yield*

$$0 \geq \rho'_{0,s_I}(X) > (1+\theta)\mathbb{E}'_{s_I}[X] \iff (*) = \rho'_{0,s_I}(X) - (1+\theta)\mathbb{E}'_{s_I}[X] > 0. \tag{16}$$

The claim holds with strict inequality (i.e. $\frac{\partial}{\partial\beta}L_0(\theta,\beta) < 0$*) unless* $s_I(\theta,\beta) = 0$*.*

Proof (Proposition 2). *Recall that by definition, for any* $\beta \in [\underline{\beta}, 1]$

$$\theta_R(\beta) = \sup\left\{\theta \geq 0 : L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta)) \leq L_{1,\mathcal{N}}(s_N)\right\}$$

is the highest loading the insurance buyer would accept. As $\theta \mapsto L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta))$ *is increasing with* $\lim\limits_{\theta\to\infty} L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta)) = \infty$*, the supremum is attained and by continuity of* $L^{\theta,\beta}_{1,\mathcal{I}}$ *it holds that at* $(\theta_R(\beta),\beta)$ *the insurance buyer is indifferent between buying and not buying insurance, i.e.*

$$\begin{aligned} L^{\theta_R(\beta),\beta}_{1,\mathcal{I}}(s_I(\theta_R(\beta),\beta)) &= L_{\mathcal{N}}(s_N) \\ (1+\theta_R(\beta))\,\mathbb{E}_{s_I(\theta_R(\beta),\beta)}[X] + \beta c(s_I(\theta_R(\beta),\beta)) &= \rho_{1,s_N}(X) + \beta_o c(s_N) \\ \iff (1+\theta_R(\beta))\,\mathbb{E}_{s_I(\theta_R(\beta),\beta)}[X] &= \rho_{1,s_N}(X) + \beta_o c(s_N) - \beta c(s_I(\theta_R(\beta),\beta)). \end{aligned} \tag{17}$$

Thus, the insurer's loss function on the boundary $(\theta_R(\beta),\beta)$ *is given by*

$$\begin{aligned} L_0(\theta_R(\beta),\beta) &= \rho_{0,s_I(\theta_R(\beta),\beta)}(X) - (1+\theta_R(\beta))\mathbb{E}_{s_I(\theta_R(\beta),\beta)}[X] + (1-\beta)c(s_I(\theta_R(\beta),\beta)) \\ &\overset{(17)}{=} \rho_{0,s_I(\theta_R(\beta),\beta)}(X) + c(s_I(\theta_R(\beta),\beta)) - \rho_{1,s_N}(X) - \beta_o c(s_N) < 0. \end{aligned}$$

The total derivative of $L_0(\theta_R(\beta),\beta)$ *w.r.t.* β *at* $s = s_I(\theta_R(\beta),\beta)$ *is given by*

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}\beta}L_0(\theta_R(\beta),\beta) &= \frac{\mathrm{d}}{\mathrm{d}\beta}\rho_{0,s_I(\theta_R(\beta),\beta)}(X) + \frac{\mathrm{d}}{\mathrm{d}\beta}c(s_I(\theta_R(\beta),\beta)) + \underbrace{\frac{\mathrm{d}}{\mathrm{d}\beta}\Big(- \rho_{1,s_N}(X) - \beta_o c(s_N)\Big)}_{=0} \\ &= \rho'_{0,s}(X)|_{s=s_I(\theta_R(\beta),\beta)} \left[\frac{\partial}{\partial\theta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)} \; \frac{\mathrm{d}}{\mathrm{d}\beta}\theta_R(\beta) + \frac{\partial}{\partial\beta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)}\right] \\ &\quad + c'(s)|_{s=s_I(\theta_R(\beta),\beta)}\left[\frac{\partial}{\partial\theta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)} \; \frac{\mathrm{d}}{\mathrm{d}\beta}\theta_R(\beta) + \frac{\partial}{\partial\beta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)}\right] \\ &= \Big[\underbrace{\rho'_{0,s}(X)|_{s=s_I}}_{\leq 0} + \underbrace{c'(s)|_{s=s_I}}_{\geq 0}\Big] \underbrace{\Big[\underbrace{\frac{\partial}{\partial\theta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)}}_{\geq 0 \text{ by Cor. 2}} \; \underbrace{\frac{\mathrm{d}}{\mathrm{d}\beta}\theta_R(\beta)}_{\leq 0 \text{ by Cor. 3}} + \underbrace{\frac{\partial}{\partial\beta}s_I(\theta,\beta)|_{\theta=\theta_R(\beta)}}_{\leq 0 \text{ by Cor. 2}}\Big]}_{\leq 0}. \end{aligned} \tag{18}$$

While it follows from the previous computations that the second factor is non-positive, the sign of the first factor is not yet determined, as the derivatives of the risk measure and the cost w.r.t. s *are of opposite signs. However, we show that their sum is always positive when evaluated at the optimal solution* $s_I(\theta_R(\beta),\beta)$ *of the buyer, such that it follows immediately that the overall product in* (18) *and thus the sign of* $\frac{\mathrm{d}}{\mathrm{d}\beta}L_0(\theta_R(\beta),\beta)$ *is non-positive.*
Recall from the proof of Corollary 2 that $s_I(\theta,\beta)$ *is the global minimizer of the strictly convex function* $L^{\theta,\beta}_{1,\mathcal{I}}(s)$ *and that for any* β *there exists* $\theta_I(\beta)$ *such that*

$$\begin{aligned} &\theta \leq \theta_I(\beta) \implies [L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=0} \geq 0, \; s_I(\theta,\beta) = 0, \\ &\theta > \theta_I(\beta) \implies [L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=0} < 0, \; s_I(\theta,\beta) > 0, \; [L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=s_I} = 0. \end{aligned}$$

In the first case, the loss function is non-decreasing at $s = 0$ *and therefore by convexity non-decreasing everywhere, thus the global minimizer given by* $s_I = 0$*. In the second case, the loss function is decreasing at* 0*, thus the global minimizer an interior point on* $(0,\infty)$ *characterized by the first-order optimality condition.*
In any case, the derivative of the loss function evaluated at the global minimizer is non-negative, i.e.

$$[L^{\theta_R(\beta),\beta}_{1,\mathcal{I}}(s)]'|_{s=s_I(\theta_R(\beta),\beta)} \geq 0. \tag{19}$$

Therefore, for the first factor in (18)*, it follows*

$$\begin{aligned} &\rho'_{0,s}(X)|_{s=s_I(\theta_R(\beta),\beta)} + c'(s)|_{s=s_I(\theta_R(\beta),\beta)} \overset{\beta\leq 1}{\geq} \rho'_{0,s}(X)|_{s=s_I(\theta_R(\beta),\beta)} + \beta c'(s)|_{s=s_I(\theta_R(\beta),\beta)} \\ &\overset{(16)}{>} (1+\theta_R(\beta))\mathbb{E}'_s[X]|_{s=s_I(\theta_R(\beta),\beta)} + \beta c'(s)|_{s=s_I(\theta_R(\beta),\beta)} = \big[L^{\theta_R(\beta),\beta}_{1,\mathcal{I}}(s)\big]'|_{s=s_I(\theta_R(\beta),\beta)} \overset{(19)}{\geq} 0, \end{aligned}$$

implying the claim of the proposition and Theorem 1.
Note again that unless $\theta_R(\beta) < \theta_I(\beta)$ *implying* $s_I(\theta_R(\beta),\beta) = 0$*, the second factor of the derivative is negative (instead of non-positive), leading to uniqueness of the solution.*

A.5 Case Study: Self-Protection with a Pareto Loss

To illustrate the results for a single insurance buyer in a self-protection setting, we now consider a loss with c.d.f.

$$F_{X,s}(x) = (1 - p(s)) + p(s)F_Y(x),\ x \geq 0,$$

where $0 \leq p(s) \leq 1$ and F_Y is the c.d.f. of a Pareto-distributed r.v. $Y \sim \text{Pareto}(\hat{x}, k)$ (compare [Bensalem et al., 2020]), i.e. a *zero-inflated Pareto distribution.* Naturally, $s \mapsto p(s)$ is assumed to be decreasing (loss probability decreases as service increases) and convex (decreasing marginal impact). Additionally, assume that $\lim\limits_{s\to\infty} p(s) > 0$, i.e. the risk of a loss can never be completely eliminated.

Assume that both parties use the PH transform, where the exponents of the distortion function express that the insurer (index r_0) is less risk averse than the buyer (index r_1): $\psi_1(u) = u^{r_1},\ \psi_0(u) = u^{r_0},\ \ r_1, r_0 \in (0,1], r_0 > r_1$. Furthermore, assume for the cost of service (recall that $s \mapsto c(s)$ is assumed increasing and strictly convex with $c(0) = 0$) the functional form $c(s) = \eta s^\gamma$, $\eta > 0$, $\gamma > 1$. It follows (see [Bensalem et al., 2020]):

$$\bar{F}_{X,s}(x) = \mathbb{P}_s(X > x) = \begin{cases} p(s) & \forall\, 0 \leq x \leq \hat{x}, \\ p(s)\left(\frac{\hat{x}}{x}\right)^k, & \forall\, \hat{x} < x, \end{cases}$$

$$\mathbb{E}_s[X] = \begin{cases} p(s)\frac{\hat{x}k}{k-1}, & \text{if } k > 1, \\ \infty, & \text{else,} \end{cases} \qquad \rho_{1,s}(X) = \begin{cases} \frac{\hat{x}r_1kp(s)^{r_1}}{r_1k-1}, & \text{if } k > \frac{1}{r_1}, \\ \infty, & \text{else,} \end{cases}$$

such that from now on, we assume $k > \frac{1}{r_1}$. In this case, $G^\beta(s)$ is given by

$$G^\beta(s) = \frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]} + (\beta_o - \beta)\frac{c(s)}{\mathbb{E}_s[X]} = \frac{r_1(k-1)}{r_1k-1}p(s)^{r_1-1} + (\beta_o - \beta)\frac{\eta s^\gamma(k-1)}{\hat{x}kp(s)},$$

which is indeed a non-decreasing function of s, as $s \mapsto p(s)$ is non-increasing and $r_1 - 1 < 0$. As $\theta_0 = \frac{r_1(k-1)p(0)^{r_1-1}}{r_1k-1} - 1$, it follows:

- If $\theta \leq \frac{r_1(k-1)p(0)^{r_1-1}}{r_1k-1} - 1$, this implies $\mathcal{N} = \emptyset$, $\mathcal{I} = [0, \infty)$;
- If $\theta > \frac{r_1(k-1)p(0)^{r_1-1}}{r_1k-1} - 1$, this implies $\mathcal{N} = [0, s_B(\theta,\beta))$, $\ \mathcal{I} = [s_B(\theta,\beta), \infty)$, where by continuity $s_B(\theta,\beta)$ solves

$$(1+\theta) - \frac{r_1(k-1)}{r_1k-1}p(s_B(\theta,\beta))^{r_1-1} - \frac{(\beta_o-\beta)(k-1)\eta}{\hat{x}k}s_B(\theta,\beta)^\gamma p(s_B(\theta,\beta))^{-1} = 0.$$

For $\theta > \theta_0$, recall that on $\mathcal{N}$, the insurance buyer minimizes

$$L_{1,\mathcal{N}}(s) = \rho_{1,s}(X) + \beta_o c(s) = \frac{\hat{x}r_1kp(s)^{r_1}}{r_1k-1} + \beta_o\eta s^\gamma.$$

From now on, let us assume that $p(s) = \frac{1}{a+s} + b,\ a \geq 1,\ b > 0$ (such that $p(0) \leq 1$), which ensures convexity of the buyer's problem.[40]

[40] Here, $s \mapsto L_{1,\mathcal{N}}(s)$ is convex iff $s \mapsto p(s)^{r_1}$ is convex. This is ensured if $(r_1-1)(p'(s))^2 + p(s)p''(s) \geq 0$, which is fulfilled for logarithmic convexity of $p(s)$ (see, e.g. [Niculescu and Persson, 2018]).

The optimality criterion for the global minimizer s_N of $L_{1,\mathcal{N}}(s)$ is then given by

$$L'_{1,\mathcal{N}}(s)|_{s=s_N} \overset{!}{=} 0 \iff s_N = \Big[- \frac{\hat{x}kr_1^2}{\beta_o\eta\gamma(r_1k-1)} p(s_N)^{r_1-1}p'(s_N)\Big]^{\frac{1}{\gamma-1}},$$

which is an interior point as $L'_{1,\mathcal{N}}(s)|_{s=0^+} < 0$ and $\lim_{s\to\infty} L'_{1,\mathcal{N}}(s) = \infty$.
Analogously, on $\mathcal{I}$, the buyer minimizes

$$L^{\theta,\beta}_{1,\mathcal{I}}(s) = (1+\theta)\mathbb{E}_s[X] + \beta c(s) = (1+\theta)\frac{\hat{x}k}{k-1}p(s) + \beta\eta s^\gamma,$$

whose derivative w.r.t. s is given by $[L^{\theta,\beta}_{1,\mathcal{I}}(s)]' = (1+\theta)\frac{\hat{x}k}{k-1}p'(s) + \beta\eta\gamma s^{\gamma-1}$, implying $[L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=0} < 0 \;\forall\; \theta > 0$ such that in all cases, $s_I(\theta,\beta) > 0$ is an interior point and the results of Corollary 2 and Section 4 hold with strict inequality, in particular the insurer's solution is unique.
The optimality criterion for $s_I(\theta,\beta)$ is given by

$$[L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=s_I} \overset{!}{=} 0 \iff s_I(\theta,\beta) = \Big[- \frac{(1+\theta)\hat{x}k}{\beta\eta\gamma(k-1)} p'(s_I(\theta,\beta))\Big]^{\frac{1}{\gamma-1}}, \qquad (20)$$

implying that $\theta \mapsto s_I(\theta,\beta)$ ($\beta \mapsto s_I(\theta,\beta)$) is increasing (decreasing) (Corollary 2).
The insurer's problem (5) in this setting becomes

$$\min_{(\theta,\beta)\in[0,\theta_R(\beta)]\times[\underline{\beta},1]} \frac{\hat{x}r_0k}{r_0k-1}\Big(\frac{1}{s_I(\theta,\beta)+a}+b\Big)^{r_0} - (1+\theta)\frac{\hat{x}k}{k-1}\Big(\frac{1}{s_I(\theta,\beta)+a}+b\Big) + (1-\beta)\eta s_I(\theta,\beta)^\gamma,$$

where $s_I(\theta,\beta)$ is characterized by Equation (20). Applying the results of Section 4 yields that the insurer's optimal solution lies in the set $\{(\theta,1),\ \theta \in [0,\underline{\theta}]\}$, i.e. the insurer's problem simplifies to

$$\min_{\theta\in[0,\theta_R(1)]} L_0(\theta,1) = \frac{\hat{x}r_0k}{r_0k-1}\Big(\frac{1}{s_I(\theta,1)+a}+b\Big)^{r_0} - (1+\theta)\frac{\hat{x}k}{k-1}\Big(\frac{1}{s_I(\theta,1)+a}+b\Big),$$

which can be shown to be monotone in θ (see [Bensalem et al., 2020]). We illustrate the buyer's and insurer's solution in Figures 6 and 7, respectively, for one exemplary set of parameters.[41]

- For small θ ($0 \le \theta \le \theta_N$), i.e. the left region in Panel 6a where $s_B(\theta,1) < \min\{s_N,\ s_I(\theta,1)\}$, both s_N and $s_I(\theta,1)$ lie in $[s_B(\theta,1),\infty) = \mathcal{I}$. Thus, the solution for $L_{1,\mathcal{N}}$ is $s_B(\theta,1)$ (Corollary 2), while the problem for $L_{1,\mathcal{I}}$ has an interior solution on $\mathcal{I}$, which is also the global solution.
- For large θ, i.e. the right region in Panel 6a where $s_B(\theta,1) > s_I(\theta,1) > s_N$, both s_N and $s_I(\theta,1)$ lie in $[0, s_B(\theta,1)) = \mathcal{N}$, yielding s_N as global solution.
- The region where θ is such that $s_N < s_B(\theta,1) < s_I(\theta,1)$ is where both problems have interior solutions on their domains, thus to determine the global solution, it remains to compare the objective functions to find the boundary $\theta_R(1)$ (see Panel 6b).[42]

[41] The parameters in this example are: for the risk measures $r_1 = 0.5$, $r_0 = 0.6$ ($r_1 < r_0$ required), for the Pareto distr. $\hat{x} = 1$, $k = 2.5$ ($k > 1/r_1$ required), for the cost $\eta = 0.5$, $\gamma = 2$, $\beta_o = 1.2$, $\underline{\beta} = 0.05$, and for the loss probability $p(s) = \frac{1}{s+1/0.6} + 0.2$ implying $p(s) \in [\lim_{s\to\infty} p(s), p(0)] = [0.2, 0.8]$.

[42] One could compare the results in Figure 6 for different choices of β (which we know are not optimal to offer from the insurer's viewpoint): As expected, if β decreases, $s_I(\theta,\beta)$ increases for any θ (Corollary 2) and thus the size of the jump in s^* at $\theta_R(\beta)$ increases. Furthermore, as β decreases, $\theta_R(\beta)$ increases (Corollary 3).

- Note again that the optimal service demand within insurance increases with the loading, but jumps downward once the premium is too high for the contract to be acceptable – in other words, insurance and risk reduction are complements (Corollary 4).
- Panel 7a compares the maximum acceptable loading for the buyer (which as stated in Corollary 3 is non-increasing in $\beta \in [\underline{\beta}, 1]$) with the minimum acceptable loading for the insurer. For small β, i.e. in the gray region left of the vertical line, no mutually acceptable contract exists.
- This is also visible in Panel 7b depicting the insurer's loss for all parameter combinations in $\{(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]\} \cup \{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$, truncated at 0. The left of the dashed vertical line corresponds to $L_0(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]$. As stated in Proposition 2, $\beta \mapsto L_0(\theta_R(\beta), \beta)$ is decreasing, leaving the optimal solution to lie in $\{(\theta, 1),\ \theta \in [0, \underline{\theta}]\}$. The loss on this part of the boundary is depicted to the right of the vertical dashed line, and for the special case of a zero-inflated Pareto distribution is decreasing (see [Bensalem et al., 2020]), such that the unique optimal solution of the insurer's problem is $(\theta^*, \beta^*) = (\theta_R(1), 1) \approx (2.88, 1)$.
- The gray areas in Panel 7b where no negative insurer's loss can be attained correspond to the small values of β to the left of the vertical line on the first part of the boundary in Panel 7a (left gray area), and the small values $\theta \in [0, \theta_{\min}(1)]$ below the intersection of $\theta_{\min}(\beta)$ with the vertical part of the boundary (right gray area).

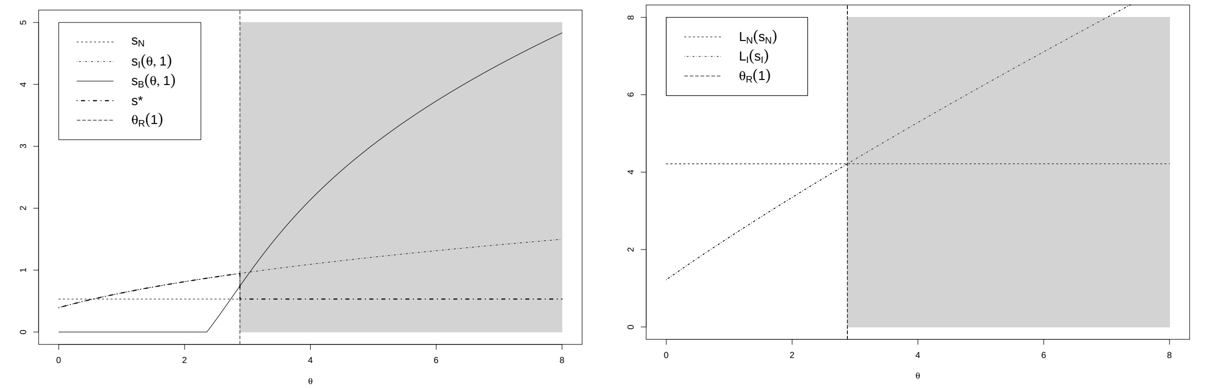

(a) $s_N, s_I(\theta, 1)$ and $s_B(\theta, 1)$.

(b) $L_{1,\mathcal{N}}(s_N)$ and $L^{\theta,1}_{1,\mathcal{I}}(s_I(\theta, 1))$.

Figure 6: Insurance buyer's solution depending on the loading θ, if he bears the full service cost ($\beta = 1$). Gray areas mark values of θ where no mutually acceptable contract exists.

A.6 Case Study: Self-Insurance with a Pareto Loss

Consider a loss with the following *zero-inflated Pareto* distribution:

$$F_{X,s}(x) = (1 - p) + pF_{Y,s}(x),\ x \geq 0,$$

where $0 < p < 1$ and $F_{Y,s}$ is the c.d.f. of a Pareto-distributed r.v. $Y \sim \text{Pareto}(\hat{x}, k(s))$, i.e. $\bar{F}_{Y,s}(x) = \left(\frac{\hat{x}}{x}\right)^{k(s)}$ for all $x > \hat{x}$ (see [Bensalem et al., 2020]).

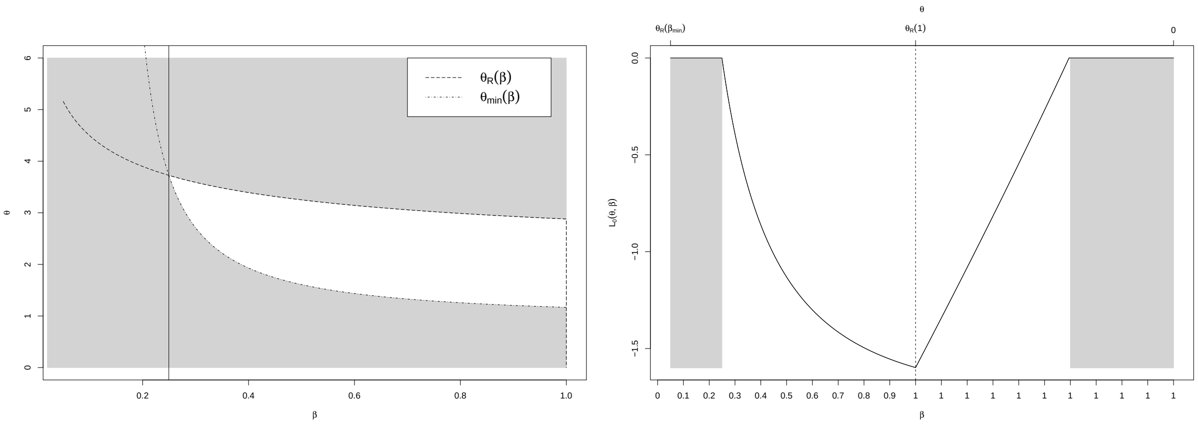

(a) $\theta_R(\beta)$ and $\theta_{\min}(\beta)$.

(b) $L_0(\theta_R(\beta),\beta),\ \beta \in [\underline{\beta},1]$ and $L_0(\theta,1),\ \theta \in [0,\underline{\theta} = \theta_R(\overline{1})]$.

Figure 7: Insurer's solution: Comparison $\theta_R(\beta)$ (maximum acceptable loading for buyer) and $\theta_{\min}(\beta)$ (minimum acceptable loading for insurer) (left panel), and insurer's loss on the boundary $\{(\theta_R(\beta),\beta),\ \beta \in [\underline{\beta},1]\} \cup \{(\theta,1),\ \theta \in [0,\underline{\theta}]\}$ (right panel). Areas where no mutually acceptable contract exists are marked gray.

This means that the service level controls the loss size via the map $s \mapsto k(s)$, which is assumed non-decreasing (service decreases loss severity), concave (decreasing marginal impact), and such that $k(0) =: z > \frac{1}{r_1} > 1$ (to ensure finiteness of $\rho_{1,s}(X)$ for any $s \geq 0$). As previously, assume $\psi_1(u) = u^{r_1},\ \psi_0(u) = u^{r_0},\ \ r_1, r_0 \in (0,1], r_0 > r_1$, and $c(s) = \eta s^{\gamma},\ \eta > 0,\ \gamma > 1$.[43] It follows (see [Bensalem et al., 2020]):

$$\bar{F}_{X,s}(x) = \begin{cases} p & \forall\ 0 \leq x \leq \hat{x}, \\ p\left(\frac{\hat{x}}{x}\right)^{k(s)} & \forall\ \hat{x} < x, \end{cases} \qquad \bar{q}_{X,s}(u) = \begin{cases} \hat{x}\left(\frac{p}{u}\right)^{1/k(s)}, & \text{if } u \in [0,p], \\ 0, & \text{if } u \in (p,1]. \end{cases}$$

$$\mathbb{E}_s[X] = \frac{p\hat{x}k(s)}{k(s)-1},\ \ k(s) > 1,\ \forall\ s \geq 0, \qquad \rho_{1,s}(X) = \begin{cases} \frac{\hat{x}r_1 k(s) p^{r_1}}{r_1 k(s) - 1}, & \text{if } k(s) > \frac{1}{r_1}, \\ \infty, & \text{else.} \end{cases}$$

Note that concavity of $s \mapsto k(s)$ implies convexity of $s \mapsto \bar{q}_{X,s}(u)$ for all $u \in (0,1)$ and that we now assume $k(s) > \frac{1}{r_1}$ for all $s \geq 0$.[44] It follows that $G^{\beta}(s)$ is given by

$$G^{\beta}(s) = \frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]} + (\beta_o - \beta)\frac{c(s)}{\mathbb{E}_s[X]} = \frac{r_1 p^{r_1 - 1}(k(s)-1)}{r_1 k(s) - 1} + (\beta_o - \beta)\frac{\eta s^{\gamma}(k(s)-1)}{\hat{x}pk(s)}.$$

Calculating the first two derivatives of the first summand w.r.t. s yields decreasingness and convexity of the ratio, i.e. $\left(\frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]}\right)' < 0,\ \ \left(\frac{\rho_{1,s}(X)}{\mathbb{E}_s[X]}\right)'' \geq 0$, implying that in this case, Assumption 3 does not hold and indeed the reverse is fulfilled (the risk measure decreases faster than the price as s increases). A sufficient condition for convexity (in s) of the second

[43] The subscripts $_0$ and $_1$ refer again to the insurer and insurance buyer, respectively.

[44] Otherwise, this would simply imply that the insurance buyer would always choose a service level of at least $s > k^{-1}(\frac{1}{r_1})$.

summand of $G^\beta(s)$ is for $c(s) = \eta s^\gamma$ and $k(s)$ to be such that for any $s \geq 0$:

$$k(s)\gamma s^{\gamma-1}k'(s) + k(s)s^\gamma k''(s) - 2s^\gamma(k'(s))^2 \geq 0. \tag{21}$$

For the remainder of this study we will assume for the service cost $\eta = 0.5$, $\gamma = 2$ (as in A.5), and for the loss severity $k(s) = \sqrt{s} + z$, where $z > \frac{1}{r_1}$. It is easily checked that Equation (21) then holds, yielding $G^\beta(s)$ as a convex function in s with $\lim\limits_{s\to\infty} G^\beta(s) = \infty$ and by continuity of the derivative decreasing for small $s > 0$, as

$$\left(G^\beta(s)\right)' = \underbrace{k'(s)}_{\xrightarrow{s\to 0}\infty}\Big[\underbrace{\frac{r_1p^{r_1-1}(r_1-1)}{(r_1k(s)-1)^2}}_{\xrightarrow{s\to 0}\text{const}<0} + \underbrace{\frac{(\beta_o-\beta)\eta s^\gamma}{p\hat{x}k(s)^2}}_{\xrightarrow{s\to 0}0}\Big] + \underbrace{\frac{(\beta_o-\beta)\eta\gamma}{p\hat{x}}\frac{k(s)-1}{k(s)}s^{\gamma-1}}_{\xrightarrow{s\to 0}0} \xrightarrow{s\to 0} -\infty.$$

Therefore, for any β, the convex map $s \mapsto G^\beta(s)$ admits a minimizer s_0 which is an interior point on $(0,\infty)$ characterized by $\big(G^\beta(s)\big)'|_{s=s_0} = 0$. The smallest loading $\theta_0(\beta)$ making $G^\beta(s)$ intersect the level $1+\theta$ for any $\theta > \theta_0$ is given by

$$\theta_0(\beta) = G^\beta(s_0) - 1,$$

such that it follows

$$\begin{aligned}
\theta \leq \theta_0(\beta) &\implies G^\beta(s) > (1+\theta)\ \forall s \geq 0,\ \mathcal{I} = [0,\infty),\ \mathcal{N} = \emptyset\\
\theta > \theta_0(\beta) &\implies \exists\ 0 \leq s_{B1}(\theta,\beta) < s_0 < s_{B2}(\theta,\beta) < \infty :\\
&\qquad \mathcal{I} = \big[0, s_{B1}(\theta,\beta)\big] \cup \big[s_{B2}(\theta,\beta),\infty\big),\ \mathcal{N} = \big(s_{B1}(\theta,\beta), s_{B2}(\theta,\beta)\big),
\end{aligned}$$

where $s_{B1}(\theta,\beta)$ and $s_{B2}(\theta,\beta)$ are defined as

$$\begin{aligned}
s_{B1}(\theta,\beta) &:= \inf\{s \in [0,\infty) : G^\beta(s) \leq 1+\theta\},\\
s_{B2}(\theta,\beta) &:= \sup\{s \in [0,\infty) : G^\beta(s) \leq 1+\theta\}.
\end{aligned}$$

Note that $s_{B1}(\theta,\beta)$ ($s_{B2}(\theta,\beta)$) is non-increasing (non-decreasing) in θ and β. The economic interpretation is straightforward: As the insurance solution gets more expensive (increase in θ or β), the interval corresponding to the no-insurance solution widens.

For $\theta > \theta_0$, recall that on $\mathcal{N} = \big(s_{B1}(\theta,\beta), s_{B2}(\theta,\beta)\big)$, the insurance buyer minimizes the loss function

$$L_{1,\mathcal{N}}(s) = \rho_{1,s}(X) + \beta_o c(s) = \frac{\hat{x}r_1p^{r_1}k(s)}{r_1k(s)-1} + \beta_o\eta s^\gamma.$$

Thus, the optimality criterion for the global minimizer s_N of $L_{1,\mathcal{N}}(s)$ is given by

$$L'_{1,\mathcal{N}}(s)|_{s=s_N} \overset{!}{=} 0 \iff s_N = \Big[\frac{\hat{x}r_1p^{r_1}k'(s_N)}{\beta_o\eta\gamma(r_1k(s_N)-1)^2}\Big]^{\frac{1}{\gamma-1}},$$

which is an interior point as $\lim\limits_{s\to0} L'_{1,\mathcal{N}}(s) = -\infty$ and $\lim\limits_{s\to\infty} L_{1,\mathcal{N}}(s) = \infty$.

Analogously, on $\mathcal{I} = [0, s_{B1}(\theta,\beta)] \cup [s_{B2}(\theta,\beta),\infty)$ the insurance buyer minimizes

$$L^{\theta,\beta}_{1,\mathcal{I}}(s) = (1+\theta)\mathbb{E}_s[X] + \beta c(s) = (1+\theta)\frac{\hat{x}pk(s)}{k(s)-1} + \beta\eta s^\gamma.$$

The optimality criterion for the global minimizer $s_I(\theta,\beta)$ is given by

$$[L^{\theta,\beta}_{1,\mathcal{I}}(s)]'|_{s=s_I} \overset{!}{=} 0 \iff s_I(\theta,\beta) = \Big[\frac{(1+\theta)\hat{x}pk'(s_I(\theta,\beta))}{\beta\eta\gamma(k(s_I(\theta,\beta))-1)^2}\Big]^{\frac{1}{\gamma-1}}. \tag{22}$$

Analogously to the self-protection case, one can argue that:

- For any β, at $\theta = 0$ it holds that $\mathcal{I} = [0, \infty)$ and $s_I(\theta, \beta)$ is also the local minimizer on $\mathcal{I}$ and the global minimizer of $L_1(\alpha, s)$ is $(\alpha^*, s^*) = (1, s_I)$.
- Define $\theta_N > 0$ as the smallest $\theta > 0$ such that $s_N \in \mathcal{N}$, i.e. $s_{B1}(\theta, \beta) < s_N$ (if $s_N \leq s_0$) or $s_N < s_{B2}(\theta, \beta)$ (if $s_N > s_0$).
- As both $s_{B2}(\theta, \beta)$ and $s_I(\theta, \beta)$ are non-decr. functions of θ, one has to distinguish the case $s_I \in \mathcal{N}$, i.e. $s_I(\theta, \beta) < s_{B2}(\theta, \beta)$ (in which automatically s_N is the global minimizer and $(\alpha^*, s^*) = (0, s_N)$ the global solution) and the case $s_I \in \mathcal{I}$, i.e. $s_{B2}(\theta, \beta) \leq s_I(\theta, \beta)$, in which one has to compare the function values $L_{1,\mathcal{I}}(s_I)$ and $L_{1,\mathcal{N}}(s_N)$ to determine the global solution.
- As $L_{1,\mathcal{N}}(s_N)$ is independent of θ, while $\theta \mapsto L_{1,\mathcal{I}}(s_I)$ is increasing, one can define the maximum feasible loading as $\theta_R(\beta) := \sup\{\theta \geq 0 : L_{1,\mathcal{I}}(s_I(\theta, \beta)) \leq L_{1,\mathcal{N}}(s_N)\}$, such that for given β for any $\theta > \theta_R(\beta)$ no insurance is preferred, i.e. $(\alpha^*, s^*) = (0, s_N)$.

The insurer's problem is therefore given by

$$\min_{(\theta,\beta)\in[0,\theta_R(\beta)]\times[\underline{\beta},1]} L_0(\theta, \beta) := \frac{\hat{x} r_0 p^{r_0} k(s_I)}{r_0 k(s_I) - 1} - (1 + \theta) \frac{\hat{x} p k(s_I)}{k(s_I) - 1} + (1 - \beta)\eta s_I^\gamma,$$

where $s_I(\theta, \beta)$ is characterized by Equation (22). In the self-insurance case, Assumption 3 does not hold anymore, and therefore neither do (necessarily) monotonicity of $\beta \mapsto L_0(\theta, \beta)$ (Proposition 1) and monotonicity of $\beta \mapsto L_0(\theta_R(\beta), \beta)$ (Proposition 2). Monotonicity (non-increasingness) of the insurer's loss in θ can be shown for the special case $\beta = 1$ (see [Bensalem et al., 2020]), but not for general $\beta \in [\underline{\beta}, 1]$. Therefore, we resort to numerical optimization (using the R package nloptr) of the program

$$\min_{\theta,\beta} L_0(\theta, \beta) \;\; s.t. \;\; L_{1,\mathcal{I}}^{\theta,\beta}(s_I(\theta, \beta)) - L_{1,\mathcal{N}}(s_N) \leq 0, \;\; 0 \leq \theta < \infty, \;\; \underline{\beta} \leq \beta \leq 1,$$

where the calculations to derive the gradients of the objective (insurer's loss function) and the constraint (insurance buyer's loss function) are given in the following subsection (for $\eta = 0.5$, $\gamma = 2$ as above). The insurance buyer's and insurer's solutions for an exemplary set of parameters are illustrated in Figures 8 and 9, respectively.[45]

- Panel 8a shows the insurance buyer's optimal service demand with and without insurance depending on the loading. As $\theta \mapsto s_I(\theta, \beta)$ is increasing, at a higher loading the insurance buyer has an incentive to purchase more service within insurance (at fixed cost). However, in this case (as remarked in [Bensalem et al., 2020] for the self-insurance case), the jump in s^* at $\theta = \theta_R(0.5)$ is positive (i.e. when switching to the no-insurance solution, more service is demanded), meaning that risk transfer demand and service demand are substitutes (contrary to Corollary 4).
- The left region of Panel 8a corresponds to the case $\theta < \theta_0$ where $\mathcal{I} = [0, \infty)$ and thus $s^* = s_I$. For $\theta \geq \theta_0$, the interval $(s_{B1}(\theta, 0.5), s_{B2}(\theta, 0.5))$ corresponding to the set $\mathcal{N}$ broadens with increasing θ.

[45] The parameters for this example are: for the risk measures $r_1 = 0.5$, $r_0 \in [r_1 + \Delta, 1]$, $\Delta = 10^{-3}$, for the service cost $\eta = 0.5$, $\gamma = 2$, $\beta_o = 1.1$, $\underline{\beta} = 0.05$, for the loss (severity) distribution $\hat{x} = 2$, $p = 0.2$, $k(s) = \sqrt{s} + z$ with $z = \frac{1}{r_1} + 0.1$. For $z \leq \frac{1}{r_1}$, one would need to calculate $\underline{s} := k^{-1}\left(\frac{1}{r_1}\right)$ and restrict the analysis to $s \in (\underline{s}, \infty)$.

- In the right region of Panel 8a, it holds $s_I(\theta, 0.5) \in \mathcal{N}$ implying $s^* = s_N$.
- For values of θ such that $s_I(\theta, 0.5) \in \mathcal{I},\ s_N \in \mathcal{N}$, one compares the insurance buyer's loss (objective function value) for both problems to determine $\theta_R(0.5)$, as illustrated in Panel 8b.
- Panel 9a compares the admissible set $(\theta, \beta) \in [0, \theta_R(\beta)] \times [\underline{\beta}, 1]$ of the insurance buyer (for $r_1 = 0.5$) with the admissible set of the insurer $(\theta, \beta) \in [\theta_{\min}(\beta), \infty) \times [\underline{\beta}, 1]$ (for selected values $r_0 \in (r_1, 1]$), i.e. a mutually acceptable contract exists for β s.t. $\theta_{\min}(\beta) \leq \theta_R(\beta)$.
- Interestingly, $\beta \mapsto \theta_{\min}(\beta)$ is not necessarily monotone decreasing anymore, and in particular if r_0 is close to r_1 (the insurer is almost as risk-averse as the buyer), the case that no acceptable contract can be found does not only occur for small values of β (as observed before), but also for very *high* values of β. As this is perhaps counter-intuitive, it merits an explanation: The more risk-averse the insurer, the more she values risk reduction by the buyer; however, the service amount $s_I(\theta, \beta)$ the buyer is willing to (optimally) purchase to achieve risk reduction with insurance decreases with β, i.e. if service becomes too expensive, the buyer may not be willing to buy as much service as required by a risk-averse insurer. This occurs due to the property of the self-insurance case that for any increase in service, the risk measure decreases faster than the price of insurance and therefore, while buying a unit of service without insurance is relatively more expensive ($\beta_o > 1$), the decrease in risk (which the buyer considers without insurance) may overcompensate this relative to the smaller decrease in price within insurance. This implies that if the risk aversions of buyer and insurer are similar, **a mutually acceptable contract can only be found if the cost of risk reduction service is shared**.
- The bold part of the boundary in Panel 9a marks the set of optimal solutions (θ^*, β^*) obtained by solving (numerically) the Stackelberg game for $r_0 \in [r_1 + \Delta, 1]$. This illustrates that contrary to the self-protection case, $\beta \mapsto L_0(\theta_R(\beta), \beta)$ is not necessarily monotone decreasing anymore.
- The solid line in Panel 9b shows the optimal share of service cost β^* burdened on the buyer depending on the absolute difference $r_0 - r_1$. The more risk-averse the insurer is relative to the buyer, the more she will incentivise risk reduction by *subsidizing* service, i.e. by optimally offering a contract with lower β. If the insurer is much less risk-averse than the buyer, she will not subsidize service any longer ($\beta^* = 1$), as the partial service cost is no longer overcompensated by her subjective gain of insuring a reduced risk.
- The dashed line in Panel 9b shows the optimally attainable value of the insurer's objective function depending on the absolute difference $r_0 - r_1$. Naturally, the insurer's obtainable gain (negative loss) decreases as she becomes more risk-averse.
- While the insurer's loss is again monotone decreasing on $\{(\theta, 1),\ \theta \in [0, \theta_R(1)]\}$ (this was shown for the special Pareto case in [Bensalem et al., 2020]), this is not necessarily the case on $\{(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]\}$ anymore. The optimal $\beta^* < 1$ and the corresponding insurer's loss are depicted in Panels 9c and 9d, for the parameter choices $r_0 = 0.6$ and $r_0 = 0.51$ (very close to r_1), respectively. Panel 9d illustrates the phenomenon observed in Panel 9a that the insurer cannot obtain a negative loss for $\beta \to 1$ (as well as for small β).

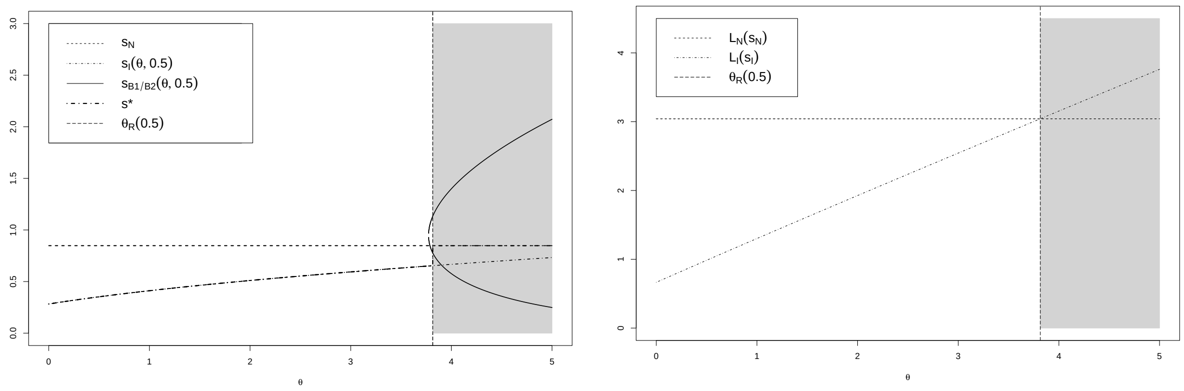

(a) $s_N, s_I(\theta, 0.5)$ and $s_{B1}(\theta, 0.5), s_{B2}(\theta, 0.5)$.

(b) $L_{1,\mathcal{N}}(s_N)$ and $L_{1,\mathcal{I}}^{\theta,0.5}(s_I(\theta, 0.5))$.

Figure 8: The figure illustrates the insurance buyer's solution depending on the loading θ, if the insurance buyer has to bear half of the service cost ($\beta = 0.5$). Gray areas again represent values of θ where no contract is closed.

Derivatives of insurer's loss in the self-insurance case

Recall that the insurer's loss function is given by

$$L_0(\theta, \beta) = \frac{\hat{x} r_0 p^{r_0} k(s_I)}{r_0 k(s_I) - 1} - (1+\theta)\frac{\hat{x} p k(s_I)}{k(s_I) - 1} + (1-\beta)c(s_I).$$

Its partial derivative w.r.t. θ is given by

$$\frac{\partial}{\partial\theta}L_0(\theta,\beta) = -\underbrace{\frac{r_0 p^{r_0}\hat{x}k'(s_I)\frac{\partial}{\partial\theta}s_I(\theta,\beta)}{(r_0k(s_I)-1)^2}}_{\leq 0} \underbrace{-p\hat{x}\frac{k(s_I)(k(s_I)-1)-(1+\theta)k'(s_I)\frac{\partial}{\partial\theta}s_I(\theta,\beta)}{(k(s_I)-1)^2}}_{<0\text{, see [Bensalem et al., 2020]}} + \underbrace{(1-\beta)c'(s_I)\frac{\partial}{\partial\theta}s_I(\theta,\beta)}_{\geq 0}, \tag{23}$$

where we use that by assumption $k'(s) \geq 0$, $k''(s) \leq 0$, and to derive $\frac{\partial}{\partial\theta}s_I(\theta,\beta)$, use that s_I is characterized by Equation (22) or equivalently (by rearranging), for the function $F_\beta(\theta, s) := \frac{s\beta(k(s)-1)^2}{p\hat{x}k'(s)} - (1+\theta)$, the tuple $(\theta, s_I(\theta,\beta))$ is a solution to $F_\beta = 0$ for any θ. Then, applying the implicit function theorem (IFT) yields (see [Bensalem et al., 2020])

$$\frac{\partial s_I(\theta,\beta)}{\partial\theta} = -\frac{\frac{\partial F_\beta(\theta,s_I)}{\partial\theta}}{\frac{\partial F_\beta(\theta,s_I)}{\partial s}} = \underbrace{\left(\frac{s_I}{1+\theta}\right)}_{\geq 0}\left(\frac{1}{1+\underbrace{\frac{2s_Ik'(s_I)}{k(s_I)-1}}_{>0}\underbrace{-\frac{k''(s_I)s_I}{k'(s_I)}}_{>0}}\right) \geq 0,$$

corroborating that $\theta \mapsto s_I(\theta,\beta)$ is non-decreasing. Equation (23) implies that in the special case $\beta = 1$ (where the third term vanishes), the insurer's loss is decreasing in θ, but this does not necessarily hold for general $\beta \in [\underline{\beta}, 1]$.

Analogously, we derive the partial derivative of the insurer's loss w.r.t. β as

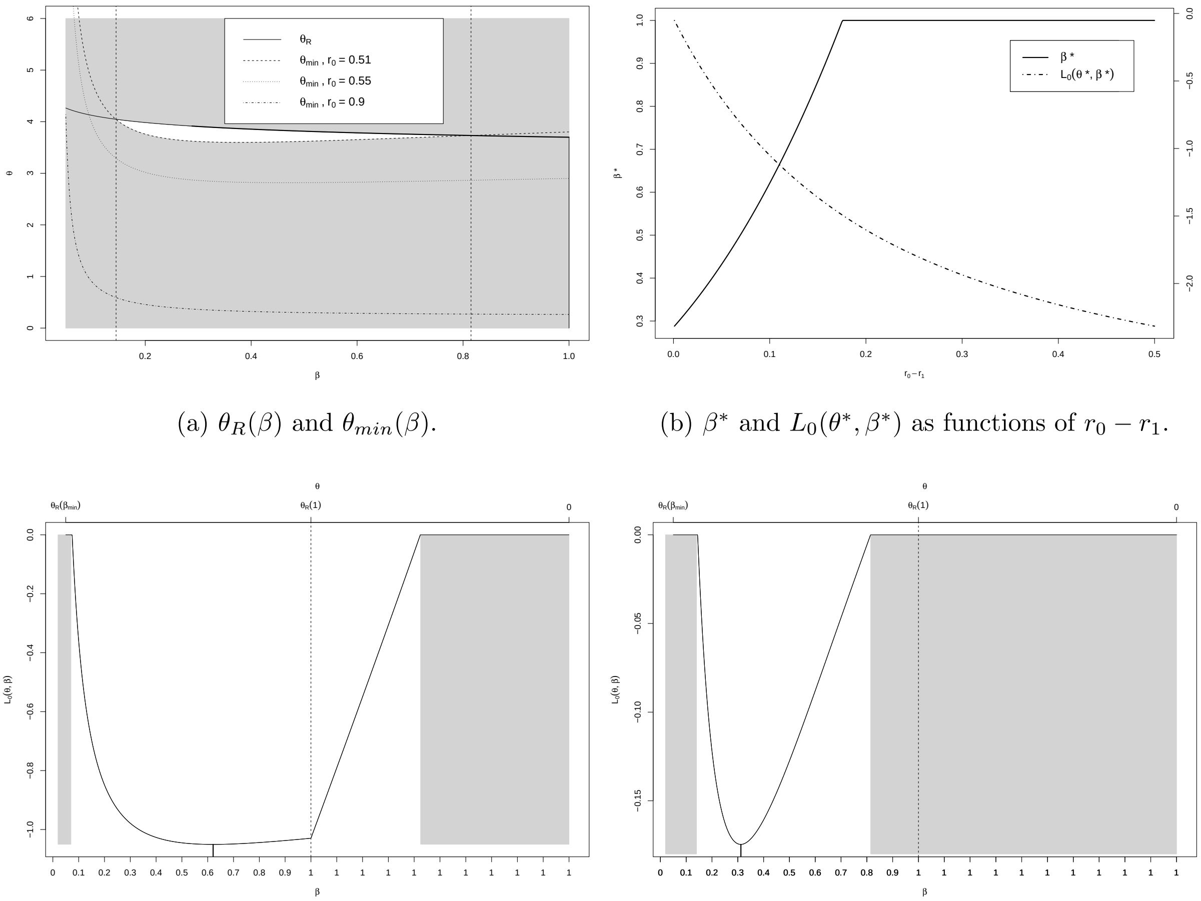

(a) $\theta_R(\beta)$ and $\theta_{min}(\beta)$.

(b) β^* and $L_0(\theta^*, \beta^*)$ as functions of $r_0 - r_1$.

(c) $r_0 = 0.6$.

(d) $r_0 = 0.51$.

Figure 9: Illustration of several aspects of the insurer's solution in the self-insurance case. Panels 9c and 9d show the insurer's loss $L_0(\theta_R(\beta), \beta),\ \beta \in [\underline{\beta}, 1]$ and $L_0(\theta, 1),\ \theta \in [0, \underline{\theta} = \theta_R(1)]$.

$$\frac{\partial}{\partial \beta} L_0(\theta, \beta) = \underbrace{-\frac{p^{r_0} r_0 \hat{x} k'(s_I) \frac{\partial}{\partial \beta} s_I(\theta, \beta)}{(r_0 k(s_I) - 1)^2}}_{\geq 0} + \underbrace{\frac{(1+\theta) p \hat{x} k'(s_I) \frac{\partial}{\partial \beta} s_I(\theta, \beta)}{(k(s_I) - 1)^2}}_{\leq 0} + \underbrace{(1-\beta) s_I \frac{\partial}{\partial \beta} s_I(\theta, \beta) - c(s_I)}_{\leq 0},$$

where, analogously to above, consider that s_I is characterized by Equation (22), or equivalently, the tuple $(\beta, s_I(\theta, \beta))$ for any β is a solution to $F_\theta(\beta, s) := \beta - (1+\theta)\frac{p\hat{x}k'(s)}{s(k(s)-1)^2} = 0$. As before, applying the IFT yields

$$\frac{\partial s_I(\theta, \beta)}{\partial \beta} = -\frac{\frac{\partial F_\theta(\beta, s_I)}{\partial \beta}}{\frac{\partial F_\theta(\beta, s_I)}{\partial s}} = \frac{1}{\beta}\left(\underbrace{\frac{k''(s_I)}{k'(s_I)}}_{\leq 0} \underbrace{-\frac{1}{s_I}}_{<0} \underbrace{-\frac{2k'(s_I)}{k(s_I) - 1}}_{\leq 0}\right)^{-1} < 0,$$

corroborating that also in this self-insurance case, $\beta \mapsto s_I(\theta, \beta)$ is non-increasing. Lastly, the gradient of the constraint $L_{1,\mathcal{I}}^{\theta,\beta}(s_I(\theta, \beta)) - L_{1,\mathcal{N}}(s_N) \leq 0$ is given by

$$
\begin{aligned}
\frac{\partial L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta))}{\partial\theta} &= p\hat{x}\frac{k(s_I)(k(s_I)-1)-(1+\theta)k'(s_I)\frac{\partial}{\partial\theta}s_I(\theta,\beta)}{(k(s_I)-1)^2} + \beta c'(s_I)\frac{\partial}{\partial\theta}s_I(\theta,\beta),\\
\frac{\partial L^{\theta,\beta}_{1,\mathcal{I}}(s_I(\theta,\beta))}{\partial\beta} &= \frac{-(1+\theta)p\hat{x}k'(s_I)\frac{\partial}{\partial\beta}s_I(\theta,\beta)}{(k(s_I)-1)^2} + \beta c'(s_I)\frac{\partial}{\partial\beta}s_I(\theta,\beta) - c(s_I).
\end{aligned}
$$

A.7 The Insurer's Problem in the Portfolio Case

A.7.1 (Directed) Loss Propagation

With the stated assumptions on $X_i,\ i\in\{1,2\}$, the portfolio loss $X := X_1 + X_2$ and its tail quantile function are given by

$$
X = \begin{cases} 0 & \text{w.p. } (1-p_1)(1-p_2),\\ L_1 & \text{w.p. } p_1(1-q)(1-p_2),\\ L_2 & \text{w.p. } (1-p_1)p_2,\\ L_1+L_2 & \text{w.p. } p_1q+p_1(1-q)p_2, \end{cases}
\quad
\overline{q}_{X,s}(u) = \begin{cases} 0 & u\in(p_1+p_2-p_1p_2,1],\\ L_1 & u\in(p_2+p_1q-p_1p_2q, p_1+p_2-p_1p_2],\\ L_2 & u\in(p_1q+p_1p_2-p_1p_2q, p_2+p_1q-p_1p_2q],\\ L_1+L_2 & u\in[0,p_1q+p_1p_2-p_1p_2q], \end{cases}
$$

where the dependence on $s_i,\ i\in\{1,2\}$, is omitted for brevity. The insurer's portfolio risk measure with $\mathbf{s} := (s_1,s_2)$ (for $\psi(u)=u^{r_0},\ r_0\in(0,1)$) is then given by

$$
\begin{aligned}
\rho_{0,\mathbf{s}}(X) &= L_1[(p_1+p_2-p_1p_2)^{r_0} - (p_2+p_1q-p_1p_2q)^{r_0}] + L_2[(p_2+p_1q-p_1p_2q)^{r_0}\\
&\quad - (p_1q+p_1p_2-p_1p_2q)^{r_0}] + (L_1+L_2)(p_1q+p_1p_2-p_1p_2q)^{r_0}\\
&= L_1[(p_1+p_2-p_1p_2)^{r_0} + (p_1q+p_1p_2-p_1p_2q)^{r_0}] + (L_2-L_1)(p_2+p_1q-p_1p_2q)^{r_0}. \qquad (24)
\end{aligned}
$$

The general difference between the single-contract and the portfolio case, i.e. that Equation (7) may hold, is illustrated in Figure 3 for an exemplary set of parameters in this case of directed loss propagation.

Calculation of gradients for sequential contract closure (SEQ 21)
Recall the insurer's objective function in Equation (10) for contract 1 is given by

$$
L^{\text{seq}}_{0,1}(\theta_1,\beta_1) = \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_{R,2}(1),1)}(X) - (1+\theta_1)\mathbb{E}_{s_{I1}(\theta_1,\beta_1)}(X_1) + (1-\beta_1)c(s_{I1}(\theta_1,\beta_1)),
$$

where the portfolio risk measure is given in (24). Its partial derivatives w.r.t. θ_1 and β_1 are thus given by

$$
\begin{aligned}
\frac{\partial L^{\text{seq}}_{0,1}(\theta_1,\beta_1)}{\partial\theta_1} &= \Big[L_1\big(r_0(p_1(s_{I1})(1-p_{I2})+p_{I2})^{r_0-1}(1-p_{I2}) + (q+p_{I2}-p_{I2}q)^{r_0}r_0p_1(s_{I1})^{r_0-1}\big)\\
&\quad + (L_2-L_1)r_0\big[p_1(s_{I1})(q-p_{I2}q)+p_{I2}\big]^{r_0-1}(q-p_{I2}q)\Big]p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1}\\
&\quad - \Big(L_1p_1(s_{I1}) + (1+\theta_1)L_1p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1}\Big) + (1-\beta_1)c'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1},\\
\frac{\partial L^{\text{seq}}_{0,1}(\theta_1,\beta_1)}{\partial\beta_1} &= \Big[L_1\big(r_0(p_1(s_{I1})(1-p_{I2})+p_{I2})^{r_0-1}(1-p_{I2}) + (q+p_{I2}-p_{I2}q)^{r_0}r_0p_1(s_{I1})^{r_0-1}\big)\\
&\quad + (L_2-L_1)r_0\big[p_1(s_{I1})(q-p_{I2}q)+p_{I2}\big]^{r_0-1}(q-p_{I2}q)\Big]p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\beta_1}\\
&\quad - (1+\theta_1)L_1p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\beta_1} - c(s_{I1}) + (1-\beta_1)c'(s_{I1})\frac{\partial s_{I1}}{\partial\beta_1},
\end{aligned}
$$

where the dependence $s_{I1}(\theta_1, \beta_1)$ is omitted for brevity and $p_{I2} := p_2(s_{I2})$. Analogously to above, to derive $\frac{\partial s_{I1}}{\partial \theta_1}$ and $\frac{\partial s_{I1}}{\partial \beta_1}$, use that s_{I1} is characterized by

$$\big[L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_1)\big]'|_{s_1=s_{I1}} \overset{!}{=} 0 \iff s_{I1}(\theta_1, \beta_1) = \Big[-\frac{(1+\theta_1)L_1}{\beta_1\eta\gamma}p_1'(s_{I1})\Big]^{\frac{1}{\gamma-1}} \overset{\eta=0.5,\gamma=2}{=} -\frac{(1+\theta_1)L_1 p'(s_{I1})}{\beta_1}. \quad (25)$$

Rearranging (25) yields that (θ_1, s_{I1}) is a solution to $F_{\beta_1}(\theta_1, s) := \frac{s\beta_1}{L_1 p_1'(s)} + (1+\theta_1) = 0$ for any θ_1 and likewise (β_1, s_{I1}) is a solution to $F_{\theta_1}(\beta_1, s) := \beta_1 + (1+\theta_1)L_1\frac{p_1'(s)}{s} = 0$ for any β_1. Applying IFT then yields

$$\frac{\partial s_{I1}(\theta_1, \beta_1)}{\partial \theta_1} = -\frac{\frac{\partial F_{\beta_1}(\theta_1,s)}{\partial \theta_1}|_{s=s_{I1}}}{\frac{\partial F_{\beta_1}(\theta_1,s)}{\partial s}|_{s=s_{I1}}} = \frac{-L_1 p_1'(s_{I1})}{\beta_1\Big(1-\frac{s_{I1}p_1''(s_{I1})}{p_1'(s_{I1})}\Big)} > 0,$$

$$\frac{\partial s_{I1}(\theta_1, \beta_1)}{\partial \beta_1} = -\frac{\frac{\partial F_{\theta_1}(\beta_1,s)}{\partial \beta_1}|_{s=s_{I1}}}{\frac{\partial F_{\theta_1}(\beta_1,s)}{\partial s}|_{s=s_{I1}}} = -\frac{s_{I1}}{\Big(p_1''(s_{I1}) - \frac{p_1'(s_{I1})}{s_{I1}}\Big)L_1(1+\theta_1)} < 0.$$

The partial derivatives of the constraint $L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1, \beta_1)) - L_{\mathcal{N},1}(s_{N1}) \leq 0$ are

$$\frac{\partial}{\partial \theta_1}L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1, \beta_1)) = L_1 p_1(s_{I1}) + (1+\theta_1)L_1 p_1'(s_{I1})\frac{\partial s_{I1}}{\partial \theta_1} + \beta_1 c'(s_{I1})\frac{\partial s_{I1}}{\partial \theta_1},$$

$$\frac{\partial}{\partial \beta_1}L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1, \beta_1)) = (1+\theta_1)L_1 p_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + c(s_{I1}) + \beta_1 c'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1}.$$

Calculation of gradients for sequential contract closure (SEQ 12)
If the contracts are closed in reverse sequential order, the insurer's objective function for contract 2 is

$$L_{0,2}^{\text{seq}}(\theta_2, \beta_2) = \rho_{0,s_{I1}(\theta_{R,1}(1),1),s_{I2}(\theta_2,\beta_2)}(X) - (1+\theta_2)\mathbb{E}_{s_{I2}(\theta_2,\beta_2)}(X_2) + (1-\beta_2)c(s_{I2}(\theta_2, \beta_2)),$$

yielding its partial derivatives w.r.t. θ_2 and β_2 as

$$\begin{aligned}\frac{\partial L_{0,2}^{\text{seq}}(\theta_2, \beta_2)}{\partial \theta_2} = {} & \Big[L_1\big(r_0(p_{I1} + (1-p_{I1})p_2(s_{I2}))^{r_0-1}(1-p_{I1}) + r_0(p_{I1}q + p_{I1}(1-q)p_2(s_{I2}))^{r_0-1}p_{I1}(1-q)\big) \\ & + (L_2 - L_1)r_0\big[p_{I1}q + (1-qp_{I1})p_2(s_{I2}))\big]^{r_0-1}(1-qp_{I1})\Big]p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \theta_2} \\ & - \Big(L_2(p_{I1}q + (1-p_{I1}q)p_2(s_{I2})) + (1+\theta_2)L_2(1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \theta_2}\Big) + (1-\beta_2)c'(s_{I2})\frac{\partial s_{I2}}{\partial \theta_2}, \\ \frac{\partial L_{0,2}^{\text{seq}}(\theta_2, \beta_2)}{\partial \beta_2} = {} & \Big[L_1\big(r_0(p_{I1} + (1-p_{I1})p_2(s_{I2})))^{r_0-1}(1-p_{I1}) + r_0(p_{I1}q + p_{I1}(1-q)p_2(s_{I2}))^{r_0-1}p_{I1}(1-q) \\ & + (L_2 - L_1)r_0\big[p_{I1}q + (1-qp_{I1})p_2(s_{I2}))\big]^{r_0-1}(1-qp_{I1})\Big]p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_2} \\ & - (1+\theta_2)L_2(1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_2} - c(s_{I2}) + (1-\beta_2)c'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_2},\end{aligned}$$

where the dependence $s_{I2}(\theta_2, \beta_2)$ is omitted, again $p_{I1} := p_1(s_{I1})$, and the partial derivatives $\frac{\partial s_{I2}}{\partial \theta_2}$ and $\frac{\partial s_{I2}}{\partial \beta_2}$ are derived analogously to above as

$$\begin{aligned}
\frac{\partial s_{I2}(\theta_2,\beta_2)}{\partial\theta_2} &= \frac{-L_2(1-p_1q)p_2'(s_{I2})}{\beta_2\Big(1-\frac{s_{I2}p_2''(s_{I2})}{p_2'(s_{I2})}\Big)} > 0,\\
\frac{\partial s_{I2}(\theta_2,\beta_2)}{\partial\beta_2} &= -\frac{s_{I2}}{\Big(p_2''(s_{I2})-\frac{p_2'(s_{I2})}{s_{I2}}\Big)L_2(1+\theta_2)(1-p_1q)} < 0.
\end{aligned}$$

The partial derivatives of the constraint $L_{\mathcal{I},2}^{\theta_2,\beta_2}(s_{I2}(\theta_2,\beta_2)) - L_{\mathcal{N},2}(s_{N2}) \leq 0$ are

$$\begin{aligned}
\frac{\partial}{\partial\theta_2}L_{\mathcal{I},2}^{\theta_2,\beta_2}(s_{I2}(\theta_2,\beta_2)) &= L_2(p_{I1}q+p_2(s_{I2})(1-p_1q)) + (1+\theta_2)L_2(1-p_1q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_2} + \beta_2c'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_2},\\
\frac{\partial}{\partial\beta_2}L_{\mathcal{I},2}^{\theta_2,\beta_2}(s_{I2}(\theta_2,\beta_2)) &= (1+\theta_2)L_2(1-p_1q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\beta_2} + c(s_{I2}) + \beta_2c'(s_{I2})\frac{\partial s_{I2}}{\partial\beta_2}.
\end{aligned}$$

Calculation of gradients for simultaneous contract closure (SIM)

Recall that if the contracts are closed simultaneously, the insurer solves the four-dimensional problem stated in Equation (11), minimizing the loss function

$$\begin{aligned}
L_0^{\text{sim}}(\theta_1,\beta_1,\theta_2,\beta_2) &= \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)}(X)\\
&\quad - (1+\theta_1)\mathbb{E}_{s_{I1}(\theta_1,\beta_1)}[X_1] - (1+\theta_2)\mathbb{E}_{s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)}[X_2]\\
&\quad + (1-\beta_1)c(s_{I1}(\theta_1,\beta_1)) + (1-\beta_2)c(s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2))\\
&=: f_1(\theta_1,\beta_1,\theta_2,\beta_2) - f_2(\theta_1,\beta_1,\theta_2,\beta_2) - f_3(\theta_1,\beta_1,\theta_2,\beta_2)\\
&\quad + f_4(\theta_1,\beta_1,\theta_2,\beta_2) + f_5(\theta_1,\beta_1,\theta_2,\beta_2),
\end{aligned}$$

on the admissible set $\mathcal{A} := [0,\theta_{R,1}(\beta_1)] \times [\underline{\beta},1] \times [0,\theta_{R,2}(\beta_2)] \times [\underline{\beta},1]$. Note that

- Due to the directed nature of loss propagation, s_{I1} does not depend on θ_2,β_2,s_{I2}, implying $\frac{\partial f_2}{\partial\theta_2} = \frac{\partial f_2}{\partial\beta_2} = \frac{\partial f_4}{\partial\theta_2} = \frac{\partial f_4}{\partial\beta_2} = 0$.
- The derivatives of the portfolio risk measure and the price of insurance and service cost for firm 2, i.e. $\frac{\partial f_1}{\partial\theta_2}$, $\frac{\partial f_1}{\partial\beta_2}$, $\frac{\partial f_3}{\partial\theta_2}$, $\frac{\partial f_3}{\partial\beta_2}$, $\frac{\partial f_5}{\partial\theta_2}$, $\frac{\partial f_5}{\partial\beta_2}$, are as in (SEQ 12).
- Analogously, the partial derivatives of the price of insurance and service cost for firm 1, i.e. $\frac{\partial f_2}{\partial\theta_1}$, $\frac{\partial f_2}{\partial\beta_1}$, $\frac{\partial f_4}{\partial\theta_1}$, $\frac{\partial f_4}{\partial\beta_1}$, are as in (SEQ 21).

The remaining derivatives w.r.t. θ_1 are given by:

$$\begin{aligned}
\frac{\partial f_1(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial\theta_1} &= L_1\Big[r_0(p_{I1}+p_{I2}-p_{I1}p_{I2})^{r_0-1}\big[p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1}(1-p_{I2}) + p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_1}(1-p_{I1})\big]\\
&\quad + r_0(p_{I1}q+p_{I1}p_{I2}-p_{I1}p_{I2}q)^{r_0-1}\big[(p_{I2}+q-p_{I2}q)p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1} + p_{I1}(1-q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_1}\big]\Big]\\
&\quad + (L_2-L_1)r_0(p_{I1}q-p_{I1}p_{I2}q+p_{I2})^{r_0-1}\Big[q(1-p_2)p_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_1}\Big],\\
\frac{\partial f_3(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial\theta_1} &= (1+\theta_2)L_2\Big((1-p_{I2})qp_1'(s_{I1})\frac{\partial s_{I1}}{\partial\theta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_1}\Big),\\
\frac{\partial f_5(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial\theta_1} &= (1-\beta_2)c'(s_{I2})\frac{\partial s_{I2}}{\partial\theta_1}.
\end{aligned}$$

To derive $\frac{\partial s_{I2}}{\partial \theta_1}$, recall that s_{I2} is characterized by $\frac{(1+\theta_2)L_2(1-qp_{I1})p_2'(s_{I2})}{\beta_2} + s_{I2} = 0$, thus for $F_{\theta_2,\beta_2}(s_1,s_2) := \frac{(1+\theta_2)L_2(1-qp_1(s_1))p_2'(s_2)}{\beta_2} + s_2$, it follows as above

$$\frac{\partial s_{I2}}{\partial s_{I1}} = -\frac{\frac{\partial F_{\theta_2,\beta_2}(s_1,s_2)}{\partial s_1}}{\frac{\partial F_{\theta_2,\beta_2}(s_1,s_2)}{\partial s_2}} = \frac{(1+\theta_2)L_2p_2'(s_{I2})qp_1'(s_{I1})}{(1+\theta_2)L_2(1-qp_{I1})p_2''(s_{I2})+\beta_2}$$

and by the chain rule $\frac{\partial s_{I2}}{\partial \theta_1} = \frac{\partial s_{I2}}{\partial s_{I1}} \cdot \frac{\partial s_{I1}}{\partial \theta_1}$ where $\frac{\partial s_{I1}}{\partial \theta_1}$ has been calculated in (SEQ 21) above. The partial derivatives w.r.t. β_1 are derived analogously as

$$\begin{aligned}
\frac{\partial f_1(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial \beta_1} &= L_1\Big[r_0(p_{I1}+p_{I2}-p_{I1}p_{I2})^{r_0-1}\big[p_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1}(1-p_{I2}) + p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}(1-p_{I1})\big] \\
&\quad + r_0(p_{I1}q + p_{I1}p_{I2} - p_{I1}p_{I2}q)^{r_0-1}\big[(p_{I2}+q-p_{I2}q)p_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + p_{I1}(1-q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}\big]\big] \\
&\quad + (L_2-L_1)r_0(p_{I1}q - p_{I1}p_{I2}q + p_{I2})^{r_0-1}\Big[q(1-p_2)p_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}\Big], \\
\frac{\partial f_3(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial \beta_1} &= (1+\theta_2)L_2\Big((1-p_{I2})qp_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}\Big), \\
\frac{\partial f_5(\theta_1,\beta_1,\theta_2,\beta_2)}{\partial \beta_1} &= (1-\beta_2)c'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1},
\end{aligned}$$

where $\frac{\partial s_{I2}}{\partial \beta_1} = \frac{\partial s_{I2}}{\partial s_{I1}} \cdot \frac{\partial s_{I1}}{\partial \beta_1}$ and $\frac{\partial s_{I1}}{\partial \beta_1}$ has been derived above.
In this case, both constraints

$$\begin{aligned}
&L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1,\beta_1)) - L_{\mathcal{N},1}(s_{N1}) \le 0, \\
&L_{\mathcal{I},2}^{\theta_1,\beta_1,\theta_2,\beta_2}(s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2)) - L_{\mathcal{N},2}(s_{N2}) \le 0,
\end{aligned}$$

have to be fulfilled simultaneously. Again $\frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1,\beta_1))}{\partial \theta_2} = \frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1,\beta_1))}{\partial \beta_2} = 0$, and $\frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1,\beta_1))}{\partial \theta_1}$, $\frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1}(\theta_1,\beta_1))}{\partial \beta_1}$ and $\frac{\partial L_{\mathcal{I},2}^{\theta_1,\beta_1,\theta_2,\beta_2}(s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2))}{\partial \theta_2}$, $\frac{\partial L_{\mathcal{I},2}^{\theta_1,\beta_1,\theta_2,\beta_2}(s_{I2}(\theta_1,\beta_1,\theta_2,\beta_2))}{\partial \beta_2}$ have been calculated above for sequential contract closure. To implement the numerical optimization routine, it remains to compute

$$\begin{aligned}
\frac{\partial L_{\mathcal{I},2}^{\theta_1,\beta_1,\theta_2,\beta_2}(s_{I2})}{\partial \theta_1} &= (1+\theta_2)L_2\Big((1-p_{I2})qp_1'(s_{I1})\frac{\partial s_{I1}}{\partial \theta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \theta_1}\Big) + \beta_2 c'(s_{I2})\frac{\partial s_{I2}}{\partial \theta_1}, \\
\frac{\partial L_{\mathcal{I},2}^{\theta_1,\beta_1,\theta_2,\beta_2}(s_{I2})}{\partial \beta_1} &= (1+\theta_2)L_2\Big((1-p_{I2})qp_1'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + (1-p_{I1}q)p_2'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}\Big) + \beta_2 c'(s_{I2})\frac{\partial s_{I2}}{\partial \beta_1}.
\end{aligned}$$

A.7.2 Cyber Events at Multiple Targets

For X_1, X_2 as above, let $Z_1 := \min\{E_1, E_{12}\}$ and $Z_2 := \min\{E_2, E_{12}\}$. Then, the portfolio loss $X := X_1 + X_2$ is described by

$$X = \begin{cases} 0 & \text{if } Z_1 > T,\ Z_2 > T \implies \mathbb{P}(X=0) = e^{-(\lambda_1+\lambda_2+\lambda_{12})} =: y_{00}, \\ L_1 & \text{if } Z_1 \le T,\ Z_2 > T \implies \mathbb{P}(X=L_1) = (1-e^{-\lambda_1})e^{-(\lambda_2+\lambda_{12})} =: y_{10}, \\ L_2 & \text{if } Z_2 \le T,\ Z_1 > T \implies \mathbb{P}(X=L_2) = (1-e^{-\lambda_2})e^{-(\lambda_1+\lambda_{12})} =: y_{01}, \\ L_1+L_2 & \text{if } Z_1 \le T,\ Z_2 \le T \implies \mathbb{P}(X=L_1+L_2) = y_{11} := 1-(y_{00}+y_{10}+y_{01}), \end{cases}$$

$$\overline{F}_X(x) = \begin{cases} 1 & x < 0, \\ 1 - y_{00} & 0 \le x < L_1, \\ 1 - (y_{00} + y_{10}) & L_1 \le x < L_2, \\ 1 - (y_{00} + y_{10} + y_{01}) & L_2 \le x < L_1 + L_2, \\ 0 & L_1 + L_2 \le x. \end{cases}$$

The risk measure of the portfolio loss is thus

$$\begin{aligned} \rho(X) &= L_1[(1 - y_{00})^r - (1 - (y_{00} + y_{10}))^r] + L_2[(1 - (y_{00} + y_{10}))^r - (1 - (y_{00} + y_{10} + y_{01}))^r] \\ &\quad + (L_1 + L_2)(1 - (y_{00} + y_{10} + y_{01}))^r \\ &= L_1[(1 - y_{00})^r + (1 - (y_{00} + y_{10} + y_{01}))^r] + (L_2 - L_1)(1 - (y_{00} + y_{10}))^r. \end{aligned}$$

Prevention of systemic events: Extended calculations

By the assumptions above, the expectation and insurer's risk measure for the contract of firm 1 are given by

$$\mathbb{E}_{s_1}[X_1] = L_1(1 - e^{-(\lambda_1 + \lambda_{12}(s_1))}), \quad \rho_{0,s_1}(X_1) = L_1(1 - e^{-(\lambda_1 + \lambda_{12}(s_1))})^{r_0},$$

while the insurer's portfolio risk measure is given by

$$\rho_{0,\mathbf{s}}(X) = L_1((1 - y_{00}(\mathbf{s}))^{r_0} + (1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}) - y_{01}(\mathbf{s}))^{r_0}) + (L_2 - L_1)(1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}))^{r_0},$$

where

$$\begin{aligned} y_{00}(\mathbf{s}) &:= e^{-(\lambda_1 + \lambda_2(s_2) + \lambda_{12}(s_1))} && \Longrightarrow \quad \frac{\partial y_{00}(\mathbf{s})}{\partial s_1} = -y_{00}(\mathbf{s})\lambda'_{12}(s_1), \\ y_{10}(\mathbf{s}) &:= (1 - e^{-\lambda_1})e^{-(\lambda_2(s_2) + \lambda_{12}(s_1))} && \Longrightarrow \quad \frac{\partial y_{10}(\mathbf{s})}{\partial s_1} = -y_{10}(\mathbf{s})\lambda'_{12}(s_1), \\ y_{01}(\mathbf{s}) &:= (1 - e^{-\lambda_2(s_2)})e^{-(\lambda_1 + \lambda_{12}(s_1))} && \Longrightarrow \quad \frac{\partial y_{01}(\mathbf{s})}{\partial s_1} = -y_{01}(\mathbf{s})\lambda'_{12}(s_1). \end{aligned}$$

The derivative of the insurer's portfolio risk measure w.r.t. s_1 is given by

$$\begin{aligned} \frac{\partial \rho_{0,\mathbf{s}}(X)}{\partial s_1} &= L_1\Big[r_0(1 - y_{00}(\mathbf{s}))^{r_0 - 1}\Big(-\frac{\partial y_{00}(\mathbf{s})}{\partial s_1}\Big) + r_0(1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}) - y_{01}(\mathbf{s}))^{r_0 - 1} \\ &\quad \Big(-\frac{\partial y_{00}(\mathbf{s}) + y_{10}(\mathbf{s}) + y_{01}(\mathbf{s})}{\partial s_1}\Big)\Big] + (L_2 - L_1)r_0(1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}))^{r_0 - 1}\Big(-\frac{\partial y_{00}(\mathbf{s}) + y_{10}(\mathbf{s})}{\partial s_1}\Big) \\ &= L_1\Big[r_0(1 - y_{00}(\mathbf{s}))^{r_0 - 1}y_{00}(\mathbf{s})\lambda'_{12}(s_1) + r_0\big(1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}) - y_{01}(\mathbf{s})\big)^{r_0 - 1} \\ &\quad (y_{00}(\mathbf{s}) + y_{10}(\mathbf{s}) + y_{01}(\mathbf{s}))\lambda'_{12}(s_1)\Big] + (L_2 - L_1)r_0(1 - y_{00}(\mathbf{s}) - y_{10}(\mathbf{s}))^{r_0 - 1}(y_{00}(\mathbf{s}) + y_{10}(\mathbf{s}))\lambda'_{12}(s_1). \end{aligned} \tag{26}$$

Using that s_{I1} is characterized by

$$[L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_1)]'|_{s_1 = s_{I1}} = (1 + \theta_1)L_1 e^{-(\lambda_1 + \lambda_{12}(s_{I1}))}\lambda'_{12}(s_{I1}) + \beta_1 c'(s_{I1}) = 0,$$

similar calculations to Subsection A.7.1 yield the derivatives of s_{I1} w.r.t. θ_1 and β_1 as

$$\frac{\partial s_{I1}}{\partial \theta_1} = -\frac{L_1 e^{-(\lambda_1 + \lambda_{12}(s_{I1}))}\lambda'_{12}(s_{I1})}{\beta_1\Big(1 + s_{I1}\lambda'_{12}(s_{I1}) - s_{I1}\frac{\lambda''_{12}(s_{I1})}{\lambda'_{12}(s_{I1})}\Big)}, \tag{27}$$

$$\frac{\partial s_{I1}}{\partial \beta_1} = \frac{s_{I1}}{(1 + \theta_1)L_1 e^{-(\lambda_1 + \lambda_{12}(s_{I1}))}\Big(\lambda'_{12}(s_{I1})^2 - \lambda''_{12}(s_{I1}) + \frac{\lambda'_{12}(s_{I1})}{s_{I1}}\Big)}, \tag{28}$$

such that the gradient of the constraint $L_{\mathcal{I},1}(s_{I1}(\theta_1,\beta_1)) - L_{\mathcal{N},1}(s_{N1}) \leq 0$ w.r.t. (θ_1,β_1) is given by

$$\begin{aligned}
\frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1})}{\partial \theta_1} &= L_1(1-e^{-(\lambda_1+\lambda_{12}(s_{I1}))}) + (1+\theta_1)L_1 e^{-(\lambda_1+\lambda_{12}(s_{I1}))}\lambda_{12}'(s_{I1})\frac{\partial s_{I1}}{\partial \theta_1} + \beta_1 c'(s_{I1})\frac{\partial s_{I1}}{\partial \theta_1}, \\
\frac{\partial L_{\mathcal{I},1}^{\theta_1,\beta_1}(s_{I1})}{\partial \beta_1} &= (1+\theta_1)L_1 e^{-(\lambda_1+\lambda_{12}(s_{I1}))}\lambda_{12}'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1} + c(s_{I1}) + \beta_1 c'(s_{I1})\frac{\partial s_{I1}}{\partial \beta_1}.
\end{aligned}$$

In the case of sequential contract closure, the insurer's objective function is

$$L_{0,1}^{\text{seq}}(\theta_1,\beta_1) = \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_{R,2}(1),1)} - (1+\theta_1)\mathbb{E}_{s_{I1}(\theta_1,\beta_1)}(X_1) + (1-\beta_1)c(s_{I1(\theta_1,\beta_1)}),$$

where the derivatives of the portfolio risk measure w.r.t. θ_1 and β_1 are

$$\begin{aligned}
\frac{\partial \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_{R,2}(1),1)}}{\partial \theta_1} &= \underbrace{\frac{\partial \rho_{0,\mathbf{s}}(X)}{\partial s_1}}_{(26)}|_{s_1=s_{I1}} \underbrace{\frac{\partial s_{I1}(\theta_1,\beta_1)}{\partial \theta_1}}_{(27)}, \\
\frac{\partial \rho_{0,s_{I1}(\theta_1,\beta_1),s_{I2}(\theta_{R,2}(1),1)}}{\partial \beta_1} &= \underbrace{\frac{\partial \rho_{0,\mathbf{s}}(X)}{\partial s_1}}_{(26)}|_{s_1=s_{I1}} \underbrace{\frac{\partial s_{I1}(\theta_1,\beta_1)}{\partial \beta_1}}_{(28)}.
\end{aligned}$$

A.7.3 Cyber Events at Multiple Targets: A Multivariate Example

We now generalize the idea behind the bivariate example in Section 5.2 to a larger portfolio. While the qualitative differences to the univariate case regarding cost-sharing can already be observed in the bivariate case, this portfolio treatment gives an indication of how to generalise the underlying idea for a specific assumption of dependence through common cyber events. We use an adapted version of the setting in [Zeller and Scherer, 2021] as follows:

- **Existing portfolio:** Assume the existing portfolio consists of $(N-1) \geq 1$ homogeneous companies indexed $j \in \{1,\dots,N-1\}$, whose contracts were priced on an individual basis previously with some risk loading $\theta_{fix} > 0$ (and implicitly $\beta = 1$). One important characteristic of each company is its IT security level, denoted here by $\ell_j \in [0,1]$, which encodes the company's ability to withstand systemic attacks (see below). We assume a homogeneous level of $\ell_j = \ell_{fix} \in [0,1)$, $j \in \{1,\dots,N-1\}$, for the existing portfolio.

- **Arrival of cyber incidents and events:** Cyber incidents at each company stem from two independent Poisson arrival processes, namely from idiosyncratic incidents (independently from other companies) and systemic events (where multiple companies are affected jointly). We assume that an arrival from a systemic event can affect each company in the portfolio with equal probability $p_{syst} \in [0,1]$. Denote the random subset of the portfolio affected by an event as $\mathcal{S} \subseteq \{1,\dots,N-1\} \cup \{N\}$. Each event arrival is furthermore equipped with a mark $m \sim Unif([0,1])$ encoding the strength of the attack, where a company which is affected by an event suffers a loss iff the strength of the attack exceeds the company's security level, i.e. $\ell_j < m$. In summary, for each company one can write the number of cyber losses in the period $[0, T]$ (in the following w.l.o.g. $T = 1$), denoted $N_j(T)$, as the sum of two independent Poisson r.v.s:

$$\begin{aligned}
&N_j(T) = N_j^{idio}(T) + N_j^{syst}(T),\\
&N_j^{idio}(T) \sim Poi(\lambda_j^{idio}\, T),\\
&N_j^{syst}(T) \sim Poi(\lambda^g\, T\, p_{syst}\, (1 - \ell_j)),
\end{aligned}$$

where λ_j^{idio} is the arrival rate of idiosyncratic cyber losses and λ^g is the fixed rate of the *ground process* of systemic cyber events.

As an additional company, indexed N, is to be added to the portfolio, the insurer seeks to price the new contract by choosing (θ_N, β_N). The insurer's choice will again induce a choice of service level within insurance, denoted $s_I(\theta_N, \beta_N)$, by the buyer, which will in turn affect his security level, denoted $\ell(s_I)$. Analogously to the bivariate example, we assume the following:

- **Effect of service on IT security level:** Initially, i.e. when approaching the insurer, buyer N has IT security level $\ell_0 \in [0, 1)$. By purchasing service at level $s \in [0, \infty)$, he can improve his security level according to

  $$\ell(s) = 1 - \frac{1}{\frac{1}{1-\ell_0} + s},$$

  such that $s \mapsto \ell(s)$ is increasing and concave with $\ell(0) = \ell_0$ and $\lim\limits_{s\to\infty} \ell(s) = 1$.

- **Loss probability for individual company:** By combining the above assumptions, one can use the properties of the Poisson distribution to express the loss probability for an individual company $j \in \{1, \dots, N\}$ as

  $$p(s) = \mathbb{P}(N_j(T) \geq 1) = 1 - \mathbb{P}(N_j(T) = 0) = 1 - \exp\Big(- T\big(\lambda_j^{idio} + \lambda^g\, p_{syst}\, (1 - \ell_j(s))\big)\Big), \tag{29}$$

  such that again $s \mapsto p(s)$ is decreasing.[46]

- **Mitigation of systemic events through warning mechanism:** In order to replicate the mechanism of prevention of systemic events by self-protection described in Remark 20, we assume the following: the companies affected by each common cyber event are targeted in a random order over time, represented by a random permutation of $\mathcal{S}$. The first company which is well-enough protected to withstand the attack furthermore enables the insurer to trigger a warning mechanism such that the remaining companies in the portfolio can

[46] Note that, as previously, to ensure convexity of the buyer's problem, one needs to check that furthermore the subjective loss probabilities $\psi(p(s))$ are convex in s. Furthermore, note that for consistency with the setting of this study we will approximate $\mathbb{P}(N_j(T) = 1) \approx \mathbb{P}(N_j(T) \geq 1)$, i.e. we assume that at most one loss can occur. This is reasonable in practice and for the parameters chosen below, as typically arrival rates of cyber incidents are very low, such that the probability of multiple losses at one policyholder in a single year is negligible (and may even be contractually excluded).

adjust their security in time in order to equally withstand the attack.[47] Denote the subset of the portfolio which is affected by an event after implying the warning mechanism as $\mathcal{S}^*(\ell(\mathbf{s}))$, where $\ell(\mathbf{s}) = (\ell_1, \dots, \ell_{N-1}, \ell_N) = (\ell_{fix}, \dots, \ell_{fix}, \ell(s_I))$.

- **Loss severity:** As in the self-protection case study in Appendix A.5, we assume each loss size to follow a Pareto distribution (with fixed parameters independent of s for a pure self-protection scenario).

Remark 25 (Effect of warning mechanism on almost homogeneous portfolio)**.** *Due to the homogeneity assumption for the structure of the existing portfolio, one needs to distinguish only three cases to understand the functionality of the warning mechanism:*

$$\begin{aligned} (\textit{weak attacks}) \quad & m \le \ell_{fix} \le \ell(s_I) \implies \mathcal{S}^* = \emptyset, \\ (\textit{medium attacks}) \quad & \ell_{fix} < m \le \ell(s_I) \implies \mathcal{S}^* \subset \mathcal{S}, \\ (\textit{strong attacks}) \quad & \ell_{fix} \le \ell(s_I) < m \implies \mathcal{S}^* = \mathcal{S}. \end{aligned}$$

In the case of a weak attack, the first affected firm can always withstand the attack and everyone else can be warned (which is actually unnecessary, as they could also withstand the attack at their established security level). In the case of a strong attack, every company in the targeted set suffers a loss and no warning mechanism can be triggered. In the second (and interesting) case of a medium attack, company N (who is better protected than the homogeneous rest of the portfolio) can withstand the attack and the subset of companies of $\mathcal{S}$ affected after N in the random permutation can be warned in time (essentially, their security level is artificially heightened to $\ell(s_I) > \ell_{fix}$ for this attack).

Remark 26 (Insurance Buyer's Problem)**.** *The optimization problem and its solution for buyer N are analogous to Section 3 and Appendix A.5 with loss probability $p(s)$ as in* (29)*, as he does not consider a portfolio viewpoint; compare Remark 19.*

Remark 27 (Insurer's Portfolio Loss and Optimization Problem)**.** *By the above assumptions, the insurer's portfolio loss X is given by*

$$X = \sum_{j=1}^{N} \Big(\sum_{i=1}^{N_j^{idio}(T)} L_i^{idio} + \sum_{i=1}^{N^{syst}(T)} \mathbb{1}_{j \in \mathcal{S}_i^*(\ell)} L_i^{syst} \Big) =: \sum_{j=1}^{N} X_j,$$

where $N_j^{idio}(T) \sim Poi(\lambda_j^{idio}\, T)$, $N^{syst}(T) \sim Poi(\lambda^g\, T)$, and $L_i^{idio}, L_i^{syst} \sim Pareto(\hat{x}, k)$, i.i.d. $\forall i$ are independent and $\mathcal{S}_i^*(\ell)$ is generated as described above independently for all i. Analogously to Remark 15, the insurer's problem for 'sequential' optimization of contract N is

[47] It is obvious that spreading out resp. ordering attacks from common events over time is actually more realistic than assuming strictly simultaneous losses. Admittedly, the specific warning mechanism is only realistic for certain types of cyber attacks, namely for those that do not immediately "notify" the victim that the system has been compromised or a loss typically stays undetected for some time. This may e.g. be the case for most data breaches (where affected companies may need years until they realize a leak) and malware which lingers or spreads in the system until the attacker uses an opportune moment to cause a business interruption or make a ransomware demand. In this case, only companies which are well-enough protected to realize an attack has been attempted can immediately trigger a warning, whereas compromised victims who do not recognize the breach cannot.

given by

$$\min_{(\theta_N,\beta_N)\in[0,\theta_{R,N}(\beta_N)]\times[\underline{\beta},1]} L^{seq}_{0,N}(\theta_N,\beta_N) = \rho_{0,\ell(\mathbf{s})}(X) - (1+\theta_N)\mathbb{E}_{\ell(s_I(\theta_N,\beta_N))}(X_N) + (1-\beta_N)c(s_I(\theta_N,\beta_N)) \\ - (N-1)(1+\theta_{fix})\mathbb{E}_{\ell_{fix}}(X_1), \tag{30}$$

where $\rho_{0,\ell(\mathbf{s})}(\cdot)$ denotes the insurer's risk measure, dependent on the security levels $\ell(\boldsymbol{s})$ of the portfolio. Note that the last term (premium for the existing portfolio) does not influence the optimization, but serves to check whether the solution complies with the necessary profitability condition $L^{seq}_{0,N}(\theta^*_N,\beta^*_N) < 0$. If the insurer priced the contract individually, she would solve

$$\min_{(\theta_N,\beta_N)\in[0,\theta_{R,N}(\beta_N)]\times[\underline{\beta},1]} L^{ind}_{0,N}(\theta_N,\beta_N) = \rho_{0,\ell(s_I)}(X_N) - (1+\theta_N)\mathbb{E}_{\ell(s_I(\theta_N,\beta_N))}(X_N) + (1-\beta_N)c(s_I(\theta_N,\beta_N)), \tag{31}$$

yielding $(\theta^*_N,\beta^*_N) = (\theta_R(1),1)$; see Section 4 and Appendix A.5.

Example 9. *In Figure 10 and Table 5, we report insightful aspects of the insurer's optimal solution depending on two variables:*

- *Figure 10 shows the optimal share of service cost and the resulting portfolio risk measure for varying p_{syst}, whereby λ^g is adjusted according to $\lambda^g = \frac{\lambda^g_0}{p_{syst}}$ for some $\lambda^g_0 > 0$ constant. This means that the overall expected number of cyber losses from systemic events stays constant, but for larger p_{syst} there are fewer events which affect on average more companies (as opposed to more events to on average smaller subsets for small p_{syst}).*
- *Table 5 shows the optimal parameters (θ^*_N,β^*_N) and the resulting service demand s_I and security level $\ell(s_I)$ of company N for different portfolio sizes N as well as the insurer's resulting total loss and portfolio risk measure.*

It is intuitive from the above construction that the effect of the warning mechanism (and therefore the benefit from company N having higher security standards) increases with both p_{syst} (influencing the expected size of the affected subset for each event) and the portfolio size N. This increases the insurer's willingness to subsidize service by lowering β_N for increasing p_{syst} or increasing N. Precisely this effect is visible in Panel 10a and Table 5. As before, with decreasing β^*_N, the attainable feasible risk loading θ^*_N increases (see Table 5). The higher the subsidy on service cost, the more service company N will purchase within insurance (see Corollary 2), leading to increased security of company N (and therefore a stronger effect of the warning mechanism). Furthermore, the difference in the insurer's portfolio risk between pricing contract N using a portfolio viewpoint and pricing contract N individually again increases (see Panel 10b and the lower part of Table 5). This emphasizes the importance of using a portfolio viewpoint when pricing the additional contract, in particular for large portfolios or portfolios with strong dependence from systemic events.

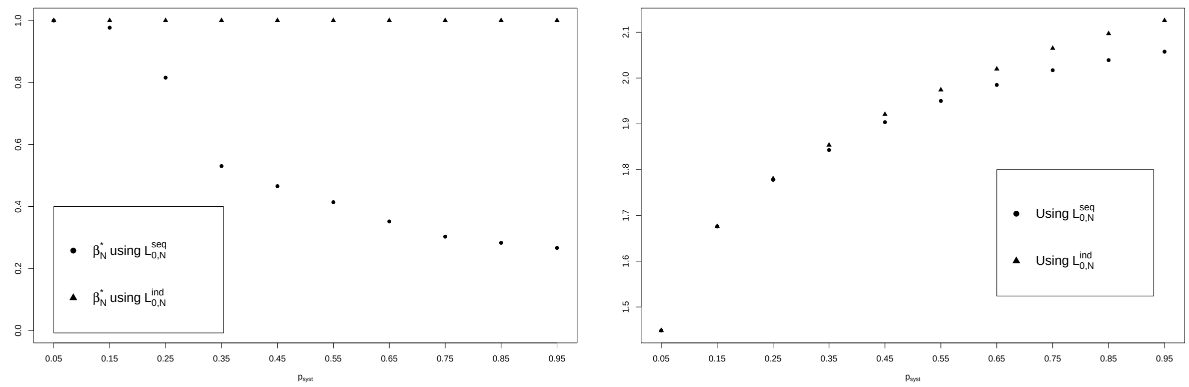

(a) β_N^* as function of p_{syst}.

(b) Insurer's optimal portfolio risk $\rho_{0,\ell(\mathbf{s})}(X)$.

Figure 10: Aspects of the insurer's solution in the portfolio case with common cyber events for varying $p_{syst} \in [0.05, 0.95]$ in steps of $\Delta = 0.1$. The parameters for this example are: portfolio size $N = 10$, loss severity parameters $\hat{x} = 1, k = 10$, loss probability parameters $\ell_{fix} = \ell_0 = 0.2, \lambda^{idio} = \frac{\lambda_0^g}{1-\ell_{fix}} = 0.05$, risk aversions $r_0 = 0.7, r_1 = 0.6$, cost parameters $\eta = 0.5, \gamma = 2, \beta_o = 1.1$. The portfolio loss distribution used to approximate the insurer's portfolio risk measure (by numerical integration) is simulated based on 1.000.000 runs.

	No insurance	Insurance (using (31))	Insurance (using (30))				
			$N = 5$	$N = 10$	$N = 25$	$N = 50$	$N = 100$
β_N^*	1.1	1	0.7350	0.4504	0.2076	0.1070	0.0623
θ_N^*	0	1.8055	1.8217	1.8566	1.9408	2.0464	2.1562
s^*	0.0561	0.0971	0.1273	0.1922	0.3501	0.5554	0.7823
$\ell(s^*)$	0.2343	0.2577	0.2739	0.3067	0.3750	0.4461	0.5080
Insurer's Loss (using (31))	0	–	-0.5333	-1.1819	-3.1754	-6.5735	-13.4685
Insurer's Loss (using (30))	0	–	-0.5338	-1.1875	-3.2226	-6.7477	-13.9697
Portfolio Risk (using (31))	0	–	1.0157	1.9467	4.6921	9.1921	18.0933
Portfolio Risk (using (30))	0	–	1.0118	1.9274	4.5883	8.8683	17.2908

Table 5: Aspects of the insurer's solution in the portfolio case with common cyber events with varying portfolio size N. All parameters are as in Figure 10 apart from $p_{syst} = 0.5$ fixed. Due to the (realistically) small rates $\lambda^{idio} = \frac{\lambda_0^g}{1-\ell_{fix}} = 0.05$, the absolute differences of the total loss and portfolio risk measure are only minor (particularly for small portfolio sizes). Therefore, we additionally report the results for another set of parameters (higher rate of systemic events) in Table 6. Note that the qualitative observations do not change, the effects are just more pronounced (as expected).

	No insurance	Insurance (using (31))	Insurance (using (30))				
			$N = 5$	$N = 10$	$N = 25$	$N = 50$	$N = 100$
β_N^*	1.1	1	0.4766	0.3007	0.1650	0.0919	0.0569
θ_N^*	0	0.6079	0.7135	0.8010	0.9807	0.8523	1.4374
s^*	0.1798	0.2960	0.5275	0.7347	1.0986	1.4475	2.0933
$\ell(s^*)$	0.3006	0.3532	0.4374	0.4961	0.5742	0.6293	0.7009
Insurer's Loss (using (31))	0	–	-2.2468	-5.3134	-14.4852	-29.8766	-60.8180
Insurer's Loss (using (30))	0	–	-2.2959	-5.4938	-15.2872	-31.9760	-66.4986
Portfolio Risk (using (31))	0	–	4.0259	7.9876	19.9007	39.6507	78.9922
Portfolio Risk (using (30))	0	–	3.8815	7.5812	18.5390	36.4532	71.1643

Table 6: Aspects of the insurer's solution in the portfolio case with common cyber events with varying portfolio size N. All parameters are as in Table 5 apart from $\lambda^{idio} = \frac{\lambda_0^g}{10(1-\ell_{fix})} = 0.05$, i.e. the ground process rate of systemic events is increased ten-fold. As this increases the occurrence rate of systemic events, it yields more pronounced (but qualitatively analogous) results to Table 5. We remark that we purposely did not choose this as the default example as for higher rates the quality of the approximation in (29) (see the corresponding footnote) deteriorates.

B Further Articles

B.1 Is accumulation risk in cyber systematically underestimated? [3]

Summary

This article is motivated by a phenomenon the authors have observed during exchanges with experts from the cyber insurance industry. “Cyber” has emerged as a novel type of risk and insurers have started to grasp this new market opportunity by devising cyber insurance solutions and establishing their cyber portfolios. Thus, actuaries have found themselves confronted with the challenging task of pricing these new policies and managing the risks from the resulting (rapidly growing, see, e.g., [15, 105]) portfolios. These are difficult tasks, as cyber risk possesses many characteristics which inhibit the applicability of standard actuarial approaches, e.g. the dynamically evolving threat landscape, the presence of strategic threat actors, and the difficulty of quantifying resulting economic losses, see e.g. [1] and the references therein. The two core challenges on a portfolio level, however, are the failure of the independence assumption between claim occurrences (which underlies the foundation of insurance) and the lack of suitable data on which the chosen models can be calibrated and back-tested.
The purpose of this article is to highlight that the challenge of cyber risk modelling must not be confined to the isolation of an actuarial department, and that holistic collaboration between experts along the whole insurance value chain is necessary and may help alleviate the challenges actuaries face.

The first part of the paper (Section 2) illustrates how the processes of product design, risk assessment, actuarial modelling, and claims settlement must be purposely designed and interconnected. This is the basis for meaningful data collection which enables the calibration and advancement of actuarial models. In particular, we aim at illustrating that the standard practice of data collection during claims settlement may be suitable to collect adequate data for calibrating individual loss distributions (which is sufficient in markets where claims can be assumed independent), but discards valuable information about the dependence between claim occurrences, which has dangerous implications for risk management.

The next part of the paper (Section 3.1) therefore introduces a purposely simplified version of the model for cyber incidents proposed in [1], which postulates that dependence between cyber incidents stems from the exploitation of *systemic vulnerabilities* (e.g. [29]) which may affect many companies simultaneously in a *common event* (a prominent example being the attacks on *Microsoft Exchange Server* in 2021, see, e.g., [131]).
To ensure tractability of the model and avoid the curse of dimensionality, we assume that under a “true” model (M), cyber events (to any non-empty subset of a portfolio of $K \in \mathbb{N}$ homogeneous companies) arrive according to independent, homogeneous Poisson processes, where the Poisson arrival rate only depends on the size of the subset (exchangeability). For an insurer, this would imply the necessity to estimate the vector of arrival rates $\boldsymbol{\lambda} := (\lambda^{|I|=1}, \dots, \lambda^{|I|=K})$, where $\lambda^{|I|=k}$

denotes the arrival rate of events affecting any subset of size $k \in \{1, \ldots, K\}$. As the maximum likelihood estimator of the rate of a homogeneous Poisson process is given by the sample mean over the observation period (see, e.g., [46]), an insurer would use the mean total number of observed events affecting precisely k companies based on the available data (i.e. settled claims from observed policy years) to obtain each estimator $\widehat{\lambda}^{|I|=k}$. However, as outlined above, as information about common events is not fully analyzed or even discarded, this available data may unfortunately not truthfully represent model (M), but rather a corresponding model with (partially) missing information, denoted $(\widetilde{M})$. We assume that, independently for each common event and for each company, the probability of an incident at this company being (in retrospect) correctly identified as belonging to the common event is given by $p \in [0, 1]$. Mathematically, this corresponds to thinning and superposition (see, e.g., [46]) of the Poisson arrival processes in model (M) according to weights connected to a Binomial distribution, leading to a new vector of homogeneous Poisson arrival rates $\widetilde{\boldsymbol{\lambda}}$ in model $(\widetilde{M})$. Thus, to quantify the consequences of missing information about common events, Section 3.2 compares quantities of interest between models (M) and $(\widetilde{M})$.

We first show that the marginal arrival rate (and therefore marginal loss distribution) for each company does not change when passing from (M) to $(\widetilde{M})$. This is meaningful for pricing, which is usually based on the marginal distribution only. In contrast, quantities of interest for portfolio risk measurement do change, which we show first by calculating the corresponding $\widetilde{\boldsymbol{\lambda}}$ for a given vector of rates $\boldsymbol{\lambda}$ and using the Panjer recursion algorithm (based on [112]) to derive the compound distribution of the total incident number in the portfolio over a policy year. This allows to compare the common *Value-at-Risk* measure at different levels and for different probabilities p, corroborating the intuition that high quantiles of the compound distribution are systematically lowered by missing information about common events, and that this phenomenon becomes more severe the lower the probability of correctly identifying the extent of common events.

We more generally use the connection between increasing convex order $(\geq_{icx})$ and ordering of coherent risk measures ([22]) to show that *Expected Shortfall* at any level (and any other coherent risk measure) of the total incident number in the portfolio will decrease in model $(\widetilde{M})$ compared to model (M). We use that (M) resp. $(\widetilde{M})$ correspond to collective risk models with equal expected claims amount and relatively few, large losses resp. relatively many, small losses, and provide an elementary proof that the former is larger ("more risky") in the sense of $\geq_{icx}$. We furthermore use the structure of the exchangeable model, i.e. the fact that first arrival times $(\tau_1, \ldots, \tau_K)$ of cyber incidents follow an *exchangeable Marshall-Olkin distribution* (see [98]) and that therefore bivariate survival copulas of any (τ_i, τ_j) are *Cuadras–Augé* copulas with known upper tail dependence coefficient, to compare the instantaneous joint loss arrival rates for any two companies under models (M) and $(\widetilde{M})$. We show that the joint loss arrival rate is underestimated by a factor of p^2 in model $(\widetilde{M})$, which intuitively corresponds to the probability of independently not overlooking a joint event in two companies.

We conclude by highlighting the urgent practical implications of this study for insurers.

This article contributes to the interface of cyber insurance research and practice by using a tractable mathematical model, which despite its simplicity still captures the essence of dependence between cyber incidents, to illustrate the urgent practical necessity of creating holistic cyber underwriting processes, by showing the otherwise potentially detrimental effects on portfolio risk measurement.

Individual contributions

I am the main author of this article. The idea of investigating this topic and a sketch of the results to be expected were developed jointly with my supervisor Matthias Scherer, who also made helpful suggestions regarding both content and presentation of the article during our regular discussions. I was responsible for the writing of the manuscript (the whole first draft as well as subsequent drafts based on comments by Matthias Scherer), developing the proofs of all statements contained therein, and the implementation and presentation of the numerical example.

Is accumulation risk in cyber systematically underestimated?

Gabriela Zeller* Matthias Scherer†

March 30, 2023

Abstract

Many insurers have started to underwrite cyber in recent years. In parallel, they developed their first actuarial models to cope with this new type of risk. On the portfolio level, two major challenges hereby are the adequate modelling of the dependence structure among cyber losses and the lack of suitable data based on which the model is calibrated. The purpose of this article is to highlight the importance of taking a holistic approach to cyber. In particular, we argue that actuarial modelling should not be viewed stand-alone, but rather as an integral part of an interconnected value chain with other processes such as cyber-risk assessment and cyber-claims settlement. We illustrate that otherwise, i.e. if these data-collection processes are not aligned with the actuarial (dependence) model, naïve data collection necessarily leads to a dangerous underestimation of accumulation risk. We illustrate the detrimental effects on the assessment of the dependence structure and portfolio risk by using a simple mathematical model for dependence through common vulnerabilities. The study concludes by highlighting the practical implications for insurers.

Keywords: Cyber Risk; Cyber Insurance; Accumulation Risk; Poisson process.

1 Introduction

Cyber insurance still is a relatively new, but steadily expanding market.[1] Insurers who have recently entered the market and started to establish their cyber portfolios, exploiting the ongoing growth in demand, are becoming increasingly aware of the challenges associated with

*Chair of Mathematical Finance, Technical University of Munich, Parkring 11, 85748 Garching-Hochbrück, Germany; Electronic address: gabi.zeller@tum.de; Corresponding author

†Chair of Mathematical Finance, Technical University of Munich, Parkring 11, 85748 Garching-Hochbrück, Germany; Electronic address: scherer@tum.de

[1] In 2015, the global market size was estimated at approximately $2 billion in premium, with US business accounting for around 90%. A rapid market growth was projected, with total premium reaching $20+ billion by 2025 ([7]). This estimate currently still seems within reach, with a global market size of around $7 billion in 2020 ([26]).

insuring cyber risk. These include the dynamically evolving threat landscape, interdependence of risks, heavy-tailed loss severities, and scarcity of reliable data to calibrate (nascent) actuarial models. Particularly the last point is repeated like a mantra; and indeed, while there are growing databases on cyber incidents and their consequences[2], they often do not contain the information necessary for the various tasks of an actuary. In fact, the best data source which can be adjusted to contain all details to calibrate an insurer's individual model is the insurer's own claims-settlement department. While an increasing number of claims in cyber insurance strain insurers' profitability margins, from the statistical point of view they should be welcomed as the detailed and reliable data whose lack is so frequently lamented. To make full use of the data collected in-house, however, the processes and systems around the underwriting of a cyber portfolio need to be aligned using a holistic approach, where risk assessment, product design, actuarial modelling, and claims settlement are treated as complementary activities interconnected by feedback loops.

In this article, we aim at illustrating the importance for insurers of using the current moment – namely when starting to underwrite cyber risk – to contemplate and establish data-collection processes in risk assessment and claims settlement which allow them to actually use the collected data to calibrate and refine their actuarial models continuously.
Sections 2.1 and 2.2, respectively, address the *cyber insurance value chain* in detail to illustrate the above mentioned interconnections and to introduce one particular approach to modelling dependence in cyber, namely via common vulnerabilities.
In Section 3 we introduce a (purposely simplified) mathematical model capturing such a dependence structure to illustrate that straightforward, naïve data collection necessarily leads to accumulation risk being systematically underestimated, both in the statistical and colloquial sense. We show that while this does not necessarily imply erroneous pricing of individual contracts, it may lead to a dangerous underestimation of dependence and portfolio risk. This is illustrated by comparing the common risk measures Value-at-Risk and Expected Shortfall for the total incident number in the portfolio as well as the joint loss arrival rate for any two companies in the portfolio.
Section 4 concludes and highlights the practical implications of this study for insurers.

2 Two Challenges for Cyber Insurance

2.1 A holistic approach to cyber-insurance underwriting

In recent years, various academic papers and numerous empirical studies have been devoted to proposing stochastic models for cyber risk.[3] Likewise, the establishment of cyber insurance as a new business line has occupied many insurers and industry subsidiaries such as brokers, see e.g. [6]. Whenever a new insurance line is introduced, the central tasks for actuaries will be technical pricing of the to-be-insured risks and risk management of the resulting

[2] See e.g. PRC [1] for a publicly available dataset on data breaches and e.g. the commercial providers Advisen [2] and SAS for more specialized datasets.

[3] For example, game-theoretic models based on a highly stylised understanding of the IT landscape ([17, 35]) or analyses of publicly available cyber loss data to propose frequency and severity distributions ([14, 15]), to name opposite ends of the modelling spectrum.

portfolio (or more precisely in cyber, risk management of an established portfolio, which now additionally contains risks from cyber policies). Underwriting and pricing risks can be done based on expert judgement for each risk individually[4] or – more commonly – based on a chosen mathematical model. In other words, actuaries have to devise an answer to the question: *"How (do we choose) to model cyber risk?"* Equally important, however, and often overlooked by academic papers, is the observation that it is not reasonable for actuaries to come up with a (no matter how accurate) answer to the above question in the isolation of an actuarial department. Instead, the chosen mathematical model needs to be simultaneously based on and itself be the basis of the business processes surrounding actuarial modelling along the entire economic insurance value chain. The development, calibration, and back-testing of an actuarial model are only sensible if they are based on information and data from risk assessment, product design, and claims settlement, as detailed below and illustrated in Figure 1.

- **Product design:** Before even starting to devise an actuarial model, a clear-cut definition and taxonomy of cyber risk(s) needs to be established in order to determine which aspects of *cyber* are deemed insurable (anything else should be excluded from the coverage by contract design) and which coverage components a cyber insurance policy should consist of. This product design process naturally needs to be revised regularly with the involvement of legal and market experts, as the cyber threat landscape as well as prospective clients' coverage needs evolve dynamically.
- **Risk assessment:** The risk-assessment process serves to elicit information deemed relevant to estimate a prospective policyholder's susceptibility to cyber risk. For cyber insurance, this process should naturally include an assessment of the client's IT infrastructure and existing cyber-security provisions. For an accurate assessment of such technical systems, cooperation with IT security experts is indispensable. However, how to adequately include extensive qualitative knowledge about an IT system's vulnerabilities and security into a stochastic model is a complex, unresolved issue in itself. Nevertheless, the questions asked and information gathered from prospective policyholders during the risk-assessment process should depend on the actuarial model that is subsequently used for pricing of individual contracts and risk management of the cyber portfolio.
- **Actuarial modelling:** The actuarial modelling step aims at developing a stochastic model which allows an estimation of the distribution of each policy's and the overall portfolio's loss from cyber risk. This serves as the basis for (technical) pricing and risk management. The model should be calibrated – and ideally back-tested – using adequate data (once available) and expert judgement. In summary, the choice of stochastic model depends on product design (which types of cyber losses are to be modelled) and in order to calibrate and develop it further, adequate data must be gathered through risk assessment and claims settlement.
- **Claims settlement:** Claims settlement deals with incoming claims from cyber losses in existing policies. In practice, this task is often treated completely disjoint from the above-mentioned processes (except product design), and typically conducted by legal experts

[4] This is indeed common e.g. for very large risks in industrial lines.

whose main concern is to understand the intricacies of each individual claim well enough to judge whether and to which extent it is covered by the components of the policy. The manner of data collection and storage is mostly dictated by legal (and efficiency) concerns. For cyber it is relevant to stress that technical expertise cannot be expected in a classical claims-settlement department. However, this is a crucial shortcoming: The information that needs to be collected in order to make claims data usable for model calibration is dictated by the choice of model. Vice versa, additional information collected may uncover flaws or omissions of the actuarial model and support its continuing development. Therefore, it is important to collect historical claims information with the underlying actuarial model in mind. In cyber, it is well-established consensus that any actuarial model needs to take **dependence** between cyber losses into account. The exact choice of dependence model is of course an insurer's individual decision[5], but it is clear that if one strives to calibrate such a model based on data, the model choice needs to be reflected in the data-collection process from the insurer's own claims experience.

Depending on the reader's own practical experience, interconnection of the above processes and cooperation between all stakeholders may sound like a utopia or a matter of course. We agree that for established business lines, either may be the case, depending on whether systems and processes were set up and continuously monitored intentionally or rather were allowed to *grow historically*. It is clear that as cyber insurance is just being established, now is the moment to intentionally set up this value chain in a way that enables insurers to cope with the dynamic challenges of this new and continuously evolving risk type in the future.

2.2 Dependence in cyber via common vulnerabilities

It is uncontested that a core actuarial challenge in cyber risk is the failure of the independence assumption between claim occurrences, which underlies the diversification principle in insurance. Due to increasing interconnectivity, businesses, systems, and supply chains become ever more dependent on functional IT infrastructure and crucially, more interdependent. Therefore, including the modelling of dependence in an actuarial model for cyber risk is indispensable. The actuarial literature discusses several approaches for this, most commonly using epidemic spreading on networks / graphs (e.g. [17, 35]), based on (marked / self- or cross-exciting) point processes (e.g. [10, 28, 36]), or employing copula approaches (e.g. [20, 24, 29]).

Regardless of the concrete modelling approach, dependence between cyber losses is worrisome for insurers as it may entail *accumulation risk*, which can be defined e.g. as the

> *risk of large aggregate losses from a single event or peril due to the concentration of insured risk exposed to that single event or peril.*[6]

Of course, accumulation risk as a concern is not limited to cyber insurance; other lines of business typically confronted with exposure concentrated to a single event are lines subject

[5] We will advocate for modelling common vulnerabilities as the source of dependence in cyber in the coming sections, but the exact choice of dependence modelling is irrelevant for this argument.

[6] Compare the definition of *risk exposure accumulation* by Casualty Actuarial Society (`https://www.casact.org/`).

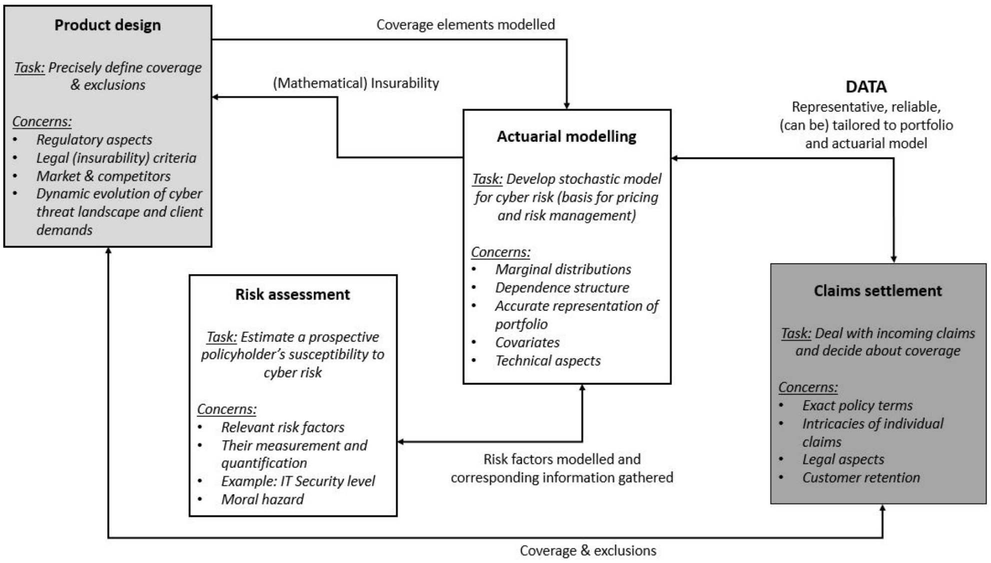
Product design
Task: Precisely define coverage & exclusions
Concerns:
Regulatory aspects
Legal (insurability) criteria
Market & competitors
Dynamic evolution of cyber threat landscape and client demands
Coverage elements modelled
(Mathematical) Insurability
Actuarial modelling
Task: Develop stochastic model for cyber risk (basis for pricing and risk management)
Concerns:
Marginal distributions
Dependence structure
Accurate representation of portfolio
Covariates
Technical aspects
DATA
Representative, reliable, (can be) tailored to portfolio and actuarial model
Risk assessment
Task: Estimate a prospective policyholder's susceptibility to cyber risk
Concerns:
Relevant risk factors
Their measurement and quantification
Example: IT Security level
Moral hazard
Risk factors modelled and corresponding information gathered
Claims settlement
Task: Deal with incoming claims and decide about coverage
Concerns:
Exact policy terms
Intricacies of individual claims
Legal aspects
Customer retention
Coverage & exclusions

Figure 1: The diagram illustrates the interconnections between different tasks in a holistic insurance value chain. While actuaries are typically mainly involved in risk assessment and actuarial modelling, there are crucial connections to other areas which must not be overlooked. In particular, the necessity to create awareness that meaningful data, which can (and should) be tailored to the chosen acturial model, is being collected daily in the claims-settlement department (usually by a completely disjoint group of experts, who do not have actuarial modelling aspects on their agenda of primary concerns) should be emphasized.

to natural catastrophes[7] or marine insurance (see e.g. [4]).

Following the classical decomposition of *risk* into a combination of threat, vulnerability, and impact (see e.g. [23]), a cyber threat only manifests itself as an incident (with potential monetary impact) if there is a corresponding vulnerability in the target system. Therefore, we postulate that any cyber incident is caused by the exploitation of a vulnerability in the company's system, where it can be distinguished between *symptomatic* and *systemic*[8] vulnerabilities (see [12, 11]), the former affecting a single company while the latter affect multiple companies simultaneously. Commonly cited examples of systemic vulnerabilities are the usage of the same operating system, cloud service provider, or payment system, affiliation with the same industry sector, or dependence on the same supplier.

Example 1. *We give two recent examples of common vulnerabilities which prominently exposed many companies to a cyber threat simultaneously. The following information and more technical details on both examples can be found in the report [33]. These examples serve to illustrate that in some cases, it might be quite obvious for an insurer to determine from incoming claims data that several cyber claims are rooted in the same common vulnerability,*

[7] For example, Hurricane Katrina has been named as the most expensive event ever to the insurance industry world-wide, see [3].

[8] We remark that some authors (see the recent survey paper [9]) employ a slightly diverging nomenclature: They denote dependency of cyber risks from common vulnerabilities as *systematic risk* and, in turn, understand *systemic risk* to mean cyber risk due to contagion effects in interconnected networks.

whereas in other cases this is very difficult to detect.

- **Microsoft Exchange:** *In the first quarter of 2021, threat actors exploited four zero-day vulnerabilities in* Microsoft Exchange Server*. The attacks drew widespread media attention due to the high number of affected companies (estimates of 60.000 victims globally, see [34]) within a short time frame, enabled by the ubiquitous use and accessibility of Exchange Servers at organizations world-wide and by their ability to be chained with other vulnerabilities. Due to the massive media coverage, leading to high awareness among companies, and the relatively clear time frame (the attacks had begun in January and were rampant during the first quarter of 2021), it was relatively easy for insurers to identify whether incoming cyber claims during (or slightly after) this time frame were rooted in one of the* Microsoft Exchange *vulnerabilities.*
- **Print Spooler / Print Nightmare:** *In the third quarter of 2021, several zero-day vulnerabilities were disclosed in* Windows Print Spooler*, another widely used service in Windows environments. As mentioned in [33], the same service was already exploited in 2010 in the so-called* Stuxnet *attacks.* Stuxnet *was a malicious worm consisting of a layered attack, where Windows systems were infected first (through zero-day vulnerabilities), but not the eventual target; i.e. the infection would have usually stayed undetected in the Windows system and seeked to propagate to certain (Siemens) PLCs (see, e.g., [18, 32]). These 2010 attacks were not immediately connected to an insurance context. However, if an analogous mechanism (e.g. through the recent* Print Spooler *vulnerabilities) were to cause cyber insurance claims, it would certainly be hard to attribute all claims to the same common vulnerability for two reasons: First, the eventual target system where the (economic) impact is caused differs from the system affected by the common vulnerability and second, the time frame is much less clear than in the previous example, as the delay between exploitation of the vulnerability and economic impact is somewhat arbitrary.*

In any case, in order to calibrate a model that uses common vulnerabilities as the source of dependence, an insurer needs to collect at least some information about the root cause for each claim to be able to estimate the dependence structure correctly. We now give a very general overview of how information on common vulnerabilities would be reflected in the insurer's risk modelling process, before introducing a more concrete, slightly simplified mathematical model in Section 3.

2.3 Notation

Assume that an insurer's portfolio consists of $K \in \mathbb{N}$ companies. From the viewpoint of each company, indexed $i \in \{1, \dots, K\}$, cyber incidents arrive according to a simple point process with corresponding counting process $(N^{(i)}(t))_{t\geq 0} = \left(|\{k \in \mathbb{N} : t_k^{(i)} \in [0,t]\}|\right)_{t\geq 0}$, in the simplest case a homogeneous Poisson process with rate $\lambda^i > 0$. This rate may differ between companies (i.e. some are assumed to be more frequently affected than others) and the main focus of cyber risk assessment (e.g. via a questionnaire, see [19] for a blueprint, or a more extensive audit for larger risks) is to gather information about characteristics which are considered relevant to determine a prospective policyholder's rate (classical covariates are e.g. company size, type and amount of data stored, types of business activities, see e.g.

[15, 30, 31]).

As the λ^i are naturally unknown, the insurer usually estimates them given past claims experience of *similar* policyholders (depending on the portfolio size, more or less homogeneous groups would be considered *similar*). The overall arrival of incoming incidents to company i is actually composed of several (assumed independent and Poisson) arrival processes (from idiosyncratic incidents and common events), i.e. the overall Poisson rate for company $i \in \{1, \dots, K\}$ decomposes into

$$\lambda^i = \lambda^{i,\text{idio}} + \sum_{s \in S_i^*} \lambda^{s,\text{syst}} > 0, \tag{1}$$

where $\lambda^{i,\text{idio}} \geq 0$ is the rate of idiosyncratic incidents arriving to company i, possibly modelled as some function of the covariates[9], $S_i^* \subseteq \{1, \dots, S\}$ is the subset of S known systemic risk factors (any common factor through which multiple companies in the portfolio could be affected simultaneously) present at company i and $\lambda^{s,\text{syst}} \geq 0$ is the overall occurrence rate of an event due to exploitation of systemic risk factor $s \in \{1, \dots, S\}$. In this modelling step, several "pitfalls" could occur:

(1) If questions about relevant covariates are omitted during risk assessment (i.e. because their influence on the frequency of cyber incidents is unknown), this may introduce a bias when estimating $\lambda^{i,\text{idio}}$ (in either direction, i.e. over-/underestimation depending on the covariates).

(2) If certain systemic risk factors are unknown and therefore not inquired about during risk assessment (e.g. no question about the choice of operating system or cloud service provider) for some or all companies, a systematic underestimation of the true rates is introduced, as the set S, resp. subsets S_i^*, do not contain all possible events.

The errors (1) and (2) should be mitigated by refining risk assessment procedures continuously based on expert input and evaluation of claims data. This leads to the main point of inquiry in this article: Given (correct) assumptions about covariates and systemic risk factors, the goal is to enable the insurer to estimate the corresponding rates, both idiosyncratic and systemic, using historical claims data. As the insurer monitors incoming claims over a policy year $[0, T]$, where typically $T = 1$, in addition to client-related data and basic claims-related data, usually a description of the incident (i.e. the order of occurrences that lead to a monetary loss) is provided by the client. This is unstructured data, and depending on the case could e.g. be given in the form of a phone conversation or e-mail report to an insurance agent or via a scanned PDF containing a report of an IT forensics expert. This information is typically reviewed by the insurance agent in order to decide whether the claim is covered, but may not or only in abbreviated form be entered into the insurer's claims database. This means that information allowing claims to be identified as stemming from the same systemic vulnerability is often not available or (fully or partly) discarded. In the following, we illustrate the detrimental effect of this omission of information about the extent of systemic events on the estimation of dependence and portfolio risk.

[9] For example, fitting a standard GLM or GAM here would be common practice.

3 Mathematical model

To quantify the effect we have introduced and discussed on a qualitative level in Section 2, we now construct a simple mathematical model which captures common events ('shocks') and allows to analyze the effect of underestimating the extent of joint events.

3.1 An exchangeable portfolio model and the modelling of missing information

We assume that the insurer's portfolio consists of $K \in \mathbb{N}$ homogeneous companies and let $\emptyset \subset I \subseteq \{1, \dots, K\}$ denote a non-empty subset of the portfolio affected by a common event. Assume that cyber events (to any set I) arrive according to independent, homogeneous Poisson processes. In theory, each subset I could potentially have a different arrival rate of common events, leading to the prohibitive complexity of needing to estimate $2^K - 1$ rates. To avoid the curse of dimensionality, we make the following assumption.

Assumption 1 (Exchangeability: Equal rates for subsets of equal size)**.** *Assume that arrival rates only depend on the number of companies in the subset, i.e. the insurer aims at estimating a vector of K arrival rates $\boldsymbol{\lambda} := (\lambda^{|I|=1}, \dots, \lambda^{|I|=K})$, where $\lambda^{|I|=k}$ denotes the arrival rate of events affecting any subset of size $k \in \{1, \dots, K\}$.*

We denote as model (M) the model given these 'true' rates $\boldsymbol{\lambda}$.[10] Assumption 1 leads to homogeneous marginal arrival rates λ^i, $i \in \{1, \dots, K\}$, for each company of

$$\lambda^i = \sum_{k=1}^{K} \frac{\lambda^{|I|=k}}{\binom{K}{k}} \binom{K-1}{k-1} = \sum_{k=1}^{K} \frac{k}{K} \lambda^{|I|=k} = \underbrace{\frac{\lambda^{|I|=1}}{K}}_{\text{idiosyncratic incidents}} + \underbrace{\sum_{k=2}^{K} \frac{k}{K} \lambda^{|I|=k}}_{\text{incidents from common events}}. \tag{2}$$

Note that (2) is a simplified formalisation of (1).
It is well-known that the maximum likelihood estimator of the rate of a homogeneous Poisson process is given by the sample mean (see e.g. [13]) over the observation period, i.e. in our case each estimator $\hat{\lambda}^{|I|=k}$ is given by the mean total number of observed events affecting precisely k companies[11], i.e. for $L > 0$ observed policy years

$$\widehat{\lambda}^{|I|=k} = \frac{1}{L} \sum_{\ell=1}^{L} \hat{n}_\ell^{|I|=k},$$

[10] Note that model (M) describes a setting where the first claim-arrival times, denoted $\tau = (\tau_1, \dots, \tau_K)$, of the companies in the portfolio follow an *exchangeable Marshall-Olkin distribution*, see [22], p. 122ff. Note that in contrast to [22], we denote by $\lambda^{|I|=k}$ the arrival rate of the Poisson process that is essentially the superimposed process of all arrival processes to subsets of size k, i.e. the rate for every particular subset of size k would be (independently of the subset) given by $\lambda_k := \frac{\lambda^{|I|=k}}{\binom{K}{k}}$. For example, for $k = 1$, $\lambda^{|I|=1}$ describes the overall rate of events affecting one single firm. As the model is exchangeable, each firm is equally likely to be affected by such an event, i.e. from the viewpoint of each of the K firms, these events arrive with rate $\lambda_1 = \frac{\lambda^{|I|=1}}{K}$.

[11] For simplicity, we assume policy years of length $T = 1$, during which the portfolio does not change.

where $\hat{n}_\ell^{|I|=k}$ is the number of observed events to subsets of size k during policy year (or simulation run) $\ell \in \{1, \dots, L\}$.

Assumption 2 (Missing information on common events)**.** *Assume that, independently for each common event to a subset of any size $|I| \geq 2$ and independently for each company in the subset, i.e. $i \in I$, the probability that the arrival at this company is correctly identified as belonging to the common event (affecting all companies in I) is given by $p \in [0, 1]$.*[12]

Example 2. *To illustrate Assumption 2, consider the following situation: A vulnerability in a commonly used software could be exploited, leading to hackers gaining access to confidential data which allowed them to defraud several companies throughout the policy year. After the policy year, when historical claims data is analyzed, all incidents in the database are first considered independent. Those incidents where detailed information is available, in this case that the original cause of the loss was the exploit of the common vulnerability, are then identified as belonging to a common event. If originally five companies were affected in this way, but only for three of them the required information was available, instead of (correctly) counting one observed event on a subset of five companies (contribution to the estimator $\widehat{\lambda}^{|I|=5}$), the insurer would (incorrectly) count one event on a subset of three companies and two independent incidents (contribution to the estimators $\widehat{\lambda}^{|I|=3}$ and twice to $\widehat{\lambda}^{|I|=1}$).*

Mathematically, Assumption 2 means that the Poisson arrival processes to subsets of size $|I| = k \geq 2$ are subject to *thinning* (with probability $(1 - p^k)$) and *superposition* of $(K - k)$ other Poisson arrival processes.

Definition 1 (Model $(\widetilde{M})$ - missing information)**.** *Assumption 2 leads to a different model, denoted $(\widetilde{M})$, with Poisson arrival rates denoted $\widetilde{\boldsymbol{\lambda}} := (\widetilde{\lambda}^{|I|=1}, \dots, \widetilde{\lambda}^{|I|=K})$ given by*

$$\widetilde{\lambda}^{|I|=1} = \lambda^{|I|=1} + \sum_{i=2}^{K} \lambda^{|I|=i} \Big[i\big(f_{Bin}(0; i, p) + f_{Bin}(1; i, p)\big) + \sum_{j=2}^{\max(i-1,2)} (i - j) f_{Bin}(j; i, p) \Big], \quad (3)$$

$$\widetilde{\lambda}^{|I|=k} = \sum_{i=k}^{K} \lambda^{|I|=i} f_{Bin}(k; i, p), \quad k \in \{2, \dots, K\}, \quad (4)$$

where $f_{Bin}(k; i, p) = \binom{i}{k} p^k (1 - p)^{i-k}$ is the p.m.f. of a Binomial distribution.

Remark 1 (Interpretation of the rates $\widetilde{\boldsymbol{\lambda}}$)**.** *The rates $\widetilde{\boldsymbol{\lambda}}$ can be interpreted as follows:*

- *For $k = K$, the rate in the model with missing information is given by*

  $$\widetilde{\lambda}^{|I|=K} = \lambda^{|I|=K} f_{Bin}(K; K, p) = \lambda^{|I|=K} p^K,$$

  i.e. the original rate thinned by the probability that all (of the K independently investigated) incidents are identified correctly. Note that for $p \in [0, 1)$, $\widetilde{\lambda}^{|I|=K} < \lambda^{|I|=K}$, i.e. the rate of events that jointly affect the whole portfolio is obviously lowered.

[12] A straightforward generalisation would be to assume different detection probabilities for different event sizes, i.e. a vector $\boldsymbol{p} := (p^{|I|=2}, \dots, p^{|I|=K})$. Intuitively, this may e.g. be used to represent the assumption that incidents from larger events are more likely to be detected, as such events are often subject to public coverage (see e.g. the *Microsoft Exchange* example above) and therefore insurers may already be alert to check if recorded claims belong to this same root cause.

- *For $1 < k < K$, the rate in the model with missing information is given by the sum of the original rate for $i = k$ thinned by the probability of classifying all k incidents correctly (summand for $i = k$) and the rates resulting from the probabilities of misclassifying events to more than k firms incorrectly such that they are counted as events to k firms (summands for $i > k$); compare Example 2. $\widetilde{\lambda}^{|I|=k}$ can thus be higher or lower than $\lambda^{|I|=k}$, depending on $\boldsymbol{\lambda}$ and p. However, in general, the cumulative rate of 'small' events (i.e. all events up to any size k) does not decrease, i.e.*

$$\sum_{i=1}^{k} \widetilde{\lambda}^{|I|=i} \geq \sum_{i=1}^{k} \lambda^{|I|=i}, \quad \forall k \in \{1, \dots, K\}.$$

- *The rate for idiosyncratic incidents in model $(\widetilde{M})$ is given by the sum of the original rate (these incidents are never "misclassified") and all the "fallout" from classifying common events incorrectly: If for an event to a subset of size i, none or only one of the firms are classified correctly, all i incidents will be counted as idiosyncratic (first part in square bracket in* (3)*); if $j \geq 2$ firms are attributed correctly, the remaining $i - j$ are classified as idiosyncratic (second part in square bracket in* (3)*). Therefore, for $p \in [0, 1)$, it holds $\widetilde{\lambda}^{|I|=1} > \lambda^{|I|=1}$, i.e. the rate of idiosyncratic incidents is increased.*

Lemma 1 (Marginal rates remain unchanged)**.** *The marginal arrival rates for each company stay unchanged between model (M) and model $(\widetilde{M})$, i.e.*

$$\widetilde{\lambda}^{i} = \lambda^{i} = \sum_{k=1}^{K} \frac{k}{K} \lambda^{|I|=k}, \quad i \in \{1, \dots, K\}.$$

Proof. Intuitively, the statement is clear, as an incorrect (non-)identification of common events does not lead to missing a claim, but to wrongly attributing its cause. A formal proof is given in Appendix A. □

The interpretation of Lemma 1 is of high practical relevance: For pricing of (cyber) insurance policies, usually only the individual loss distribution of a company is taken into account. As the marginal arrival rates stay unchanged, prices for all individual insurance contracts would stay unchanged (i.e. 'correct') between models (M) and $(\widetilde{M})$. This means that omitting information about common events would not lead to mispricing of individual policies. This identity of marginal rates is dangerous, as the crucial oversight of underestimating the extent of common events would not be evident as affecting (average) profitability, but only in a (worst-case) scenario that an unexpectedly large loss (exceeding the estimated risk measure, typically Value-at-Risk, which may be much smaller in model $(\widetilde{M})$ than the actual one in model (M), see next section) manifests.

3.2 Implications for dependence- and risk-measurement

Portfolio Value-at-Risk

Despite the marginal rates staying unchanged when moving from (M) to $(\widetilde{M})$, see Lemma 1, omitting information about common events may have dangerous implications for risk management. We first illustrate how it may lead to an underestimation of portfolio risk, measured e.g. by Value-at-Risk, denoted $\mathbf{VaR}_{1-\gamma}$, of the total incident number in the portfolio in a policy year.[13] $\mathbf{VaR}_{1-\gamma}$ for a r.v. X in an actuarial context (where positive values denote losses) is defined as

$$\mathbf{VaR}_{1-\gamma}(X) = \inf\big\{x \in \mathbb{R} : \mathbb{P}(X \leq x) \geq 1-\gamma\big\}, \quad \gamma \in (0,1). \tag{5}$$

Note that the overall incident number in a portfolio of size K follows a compound Poisson distribution, i.e.

$$S(T) := \sum_{i=1}^{N(T)} Z_i, \quad \textit{where } N(T) \sim \mathrm{Poi}\Big(T \sum_{k=1}^{K} \lambda^{|I|=k}\Big),$$
$$\{Z_i\}_{i\in\mathbb{N}} \textit{ i.i.d. with } \mathbb{P}(Z_i = k) = \frac{\lambda^{|I|=k}}{\sum_{k=1}^{K} \lambda^{|I|=k}},\ \forall k \in \{1,\dots,K\}.$$

The rate $\Big(\sum_{k=1}^{K} \lambda^{|I|=k}\Big)$ corresponds to the overall Poisson arrival rate of events (of any size), and $\{Z_i\}_{i\in\mathbb{N}}$ correspond to the associated "jump sizes" of the total incident number, i.e. the number of companies affected in the i^{th} event. Therefore, we can use the *Panjer recursion* formula (based on [27]) to compute the probability mass function (p.m.f.) and corresponding cumulative distribution function (c.d.f.) and Value-at-Risk (as in Equation (5)) of the total incident number in a policy year under models (M) and $(\widetilde{M})$ for chosen $\boldsymbol{\lambda}$ and $p \in [0,1]$. We choose an exemplary set of rates for a portfolio of size $K = 10$ as given in Table 1, where $\boldsymbol{\lambda}$ again denotes the rates of an original model (M) and $\widetilde{\boldsymbol{\lambda}}$ the rates of the corresponding model $(\widetilde{M})$ resulting from Assumption 2.

Figure 2a displays the p.m.f. under model (M) and highlights the comparison of $\mathbf{VaR}_{0.995}$ for $p = 1$ (full information, i.e. original rates), $p = 0.5$ (partial information about common events, compare Table 2), and $p = 0$ (no information about common events, i.e. complete independence assumption). Figure 2b compares $\mathbf{VaR}_{1-\gamma}$ for $(1-\gamma) \in \{0.95, 0.995\}$ and $p \in [0,1]$, based on the c.d.f. of total incident numbers under the rates $\boldsymbol{\lambda}$ and $\widetilde{\boldsymbol{\lambda}}$. This small example already highlights the importance of gathering (full!) information about the origins of cyber incidents, as otherwise the portfolio risk will be drastically underestimated.

Finally, let us mention an observation that can be made by considering the p.m.f. (and corresponding c.d.f.) for different $p \in [0,1]$, as exemplarily depicted in Figure 3: When moving from (M) to $(\widetilde{M})$, no events / incidents are missed completely, thus the c.d.f.s of the total incident number in the portfolio are **not** ordered in the sense of *usual stochastic order*, i.e. it does not hold that for all $x \geq 0 : F_{S_{\widetilde{M}}(T)}(x) \geq F_{S_M(T)}(x)$, where $S_M(T)$ (resp. $S_{\widetilde{M}}(T)$) denotes the total incident number under model (M) (resp. $(\widetilde{M})$).

We have observed, however, from the results illustrated in Table 2 and Figure 2, that this ordering of c.d.f.s does hold for certain large values of x. Figure 3b shows that indeed it holds *exactly* for large values of x, more precisely $x > x_0$ for some $x_0 \geq 0$, i.e. the so-called *single-crossing condition* or *cut-off criterion* (see e.g. [25]) is fulfilled here. This is meaningful as

[13] For the sake of simplicity, we only consider incident numbers here, as of course the results would not be qualitatively different if for an insurance application, one were to equip each incident with a (random) monetary loss size.

Model	p	$\lvert I\rvert = 1$	$\lvert I\rvert = 2$	$\lvert I\rvert = 3$	$\lvert I\rvert = 4$	$\lvert I\rvert = 5$	$\lvert I\rvert = 6$	$\lvert I\rvert = 7$	$\lvert I\rvert = 8$	$\lvert I\rvert = 9$	$\lvert I\rvert = 10$
(M)	1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
$(\widetilde{M})$	0.5	29.49	1.93	1.77	1.45	1.00	0.55	0.23	0.07	0.01	0.0010
$(\widetilde{M})$	0	55.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 1: Original rates and resulting rates for $p = 0.5$ (i.e. for each event affecting a subset of at least two firms jointly, the incident at each firm is attributed correctly to this event with probability $p = 0.5$ and otherwise incorrectly seen as independent as a result of not being able to identify the common root cause) and $p = 0$. By partially omitting information about common events, the resulting idiosyncratic rates are much increased, rates of smaller common events (here up to $|I| = 4$) are also increased, whereas rates of larger common events (here from $|I| = 6$ on) are lowered.

Model	p	λ^i ($i \in \{1, \ldots, K\}$)	$\mathbb{E}[S(T)]$	$\mathbf{VaR}_{0.95}(S(T))$	$\mathbf{VaR}_{0.99}(S(T))$	$\mathbf{VaR}_{0.995}(S(T))$
(M)	1	5.5	55	90	107	113
$(\widetilde{M})$	0.5	5.5	55	76	86	90
$(\widetilde{M})$	0	5.5	55	68	74	76

Table 2: Resulting marginal rates (homogeneous for all companies), expected total incident numbers, and risk measures $\mathbf{VaR}_{1-\gamma}(S(T))$ at three levels for $p \in \{0, 0.5, 1\}$ and $T = 1$. Crucially, marginal rates and thus expected incident numbers $\mathbb{E}[S(T)]$ do not change (by Lemma 1 and linearity), while $\mathbf{VaR}_{1-\gamma}(S(T))$ at all chosen levels is systematically lowered when common event information is partly or fully disregarded.

it is a sufficient condition for another (weaker) type of stochastic order, so-called *increasing convex order*, which has an important connection to the class of coherent risk measures; this will be addressed more generally in a subsequent section.

Quantifying dependence by joint loss arrival rate

From a practical viewpoint, the illustrations of the last section already emphasize the detrimental effects of missing information about common events. Theoretically, there are different quantities one might use to assess the extent of "missed / overlooked dependence" in model $(\widetilde{M})$ compared to the true model (M). From a risk management perspective, it is clear that simultaneous losses by multiple policyholders carry potentially greater risk than independent, diversifiable losses. Therefore, one might look at the instantaneous rate of two policyholders $i, j \in \{1, \ldots, K\},\ i \neq j$, simultaneously experiencing a cyber claim.[14] As arrivals of cyber incidents to policyholder $i \in \{1, \ldots, K\}$ follow a Poisson process with rate λ^i (see (2)), the first arrival time, denoted τ_i, follows an exponential distribution and for small $T > 0$ it holds by a first-order Taylor expansion

$$\mathbb{P}(\tau_i \leq T) = 1 - e^{-\lambda^i T} \approx 1 - (1 - \lambda^i T) = \lambda^i T \iff \frac{1}{T} \approx \frac{\lambda^i}{\mathbb{P}(\tau_i \leq T)}.$$

[14] As we are assuming an *exchangeable* model, w.l.o.g. $i = 1,\ j = 2$.

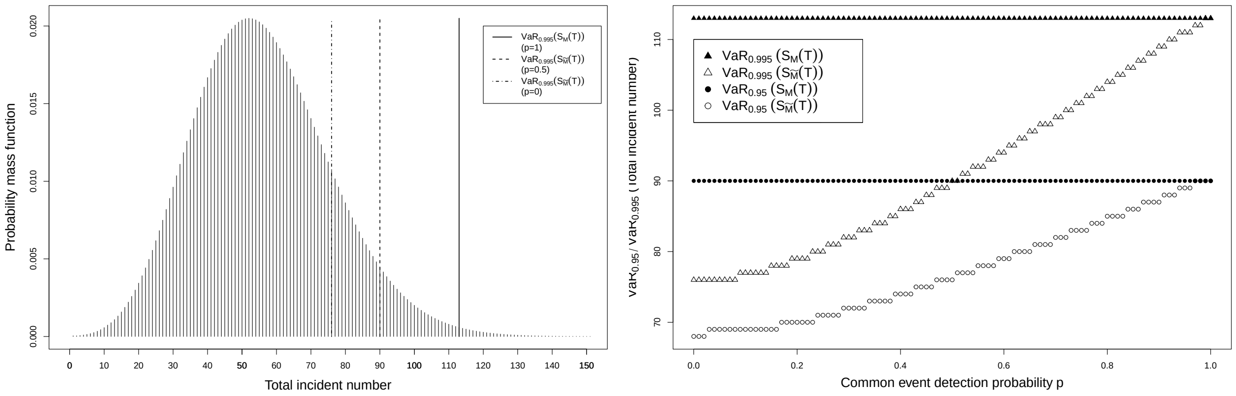

(a) $\mathbf{VaR}_{0.995}(S(T))$ for $p \in \{0, 0.5, 1\}$.

(b) $\mathbf{VaR}_{0.95/0.995}(S(T))$ for $p \in [0, 1]$.

Figure 2: Panel 2a shows the p.m.f. of the total incident number for parameters as in Table 1 and again $T = 1$. The solid vertical line depicts the corresponding $\mathbf{VaR}_{0.995}$ if full information about common events is available ($p = 1$), i.e. all incidents are classified correctly. The dashed lines depict analogously $\mathbf{VaR}_{0.995}$ for partial information ($p = 0.5$, i.e. for each event on average half of the resulting incidents are attributed correctly), and no information ($p = 0$, i.e. all incidents regarded as idiosyncratic) about common events. In both latter cases, the true risk is clearly underestimated (compare $\mathbf{VaR}_{0.995}$ for $p = 0$ with the 'true' underlying distribution!). Panel 2b shows $\mathbf{VaR}_{.}$ for $(1-\gamma) \in \{0.95, 0.995\}$ and $p \in [0, 1]$ (in steps of $\Delta = 0.01$), based on underlying rates $\boldsymbol{\lambda}$ and $\widetilde{\boldsymbol{\lambda}}$. As expected, the lower the probability p of correctly identifying a common root cause, the more severe is the resulting underestimation of the risk.

This implies for the instantaneous joint loss arrival rate

$$\lim_{T \searrow 0} \frac{\mathbb{P}(\tau_i \leq T, \tau_j \leq T)}{T} \approx \lim_{T \searrow 0} \frac{\lambda^i \mathbb{P}(\tau_i \leq T, \tau_j \leq T)}{\mathbb{P}(\tau_i \leq T)} = \lambda^i \lim_{T \searrow 0} \mathbb{P}(\tau_j \leq T \mid \tau_i \leq T) = \lambda^i \,\mathrm{LTD}_C, \tag{6}$$

where τ_i, τ_j are the first arrival times of a cyber claim to policyholders i and j, respectively, and LTD_C denotes the *lower tail dependence coefficient* of the bivariate copula C of (τ_i, τ_j). We know (see [22], p. 122ff) that by Assumption 1 the survival copula of the random vector of all K first claim-arrival times, $(\tau_1, \dots, \tau_K)$, is an *exchangeable Marshall–Olkin (eMO)* survival copula, and its two-margins (i.e. the survival copula of (τ_i, τ_j)) are bivariate *Cuadras–Augé* copulas with parameter α given by[15]:

$$\alpha = 1 - \frac{\sum_{i=1}^{K-1} \binom{K-2}{i-1} \frac{1}{\binom{K}{i}} \lambda^{|I|=i}}{\sum_{i=1}^{K} \binom{K-1}{i-1} \frac{1}{\binom{K}{i}} \lambda^{|I|=i}} = 1 - \frac{\sum_{i=1}^{K-1} \binom{K-2}{i-1} \lambda_i}{\sum_{i=1}^{K} \binom{K-1}{i-1} \lambda_i}. \tag{7}$$

From (7), some interpretation of α is immediately visible:

- Comonotonicity occurs iff only common events to the whole portfolio occur, i.e. $\alpha = 1 \iff \lambda_K > 0, \lambda_i = 0 \; \forall i \in \{1, \dots, K-1\}$;
- Independence occurs iff only idiosyncratic incidents occur, i.e. $\alpha = 0 \iff \lambda_1 > 0, \lambda_i = 0 \; \forall i \in \{2, \dots, K\}$.

[15] See the previous footnote on the relation of $\lambda^{|I|=i}$ and λ_i.

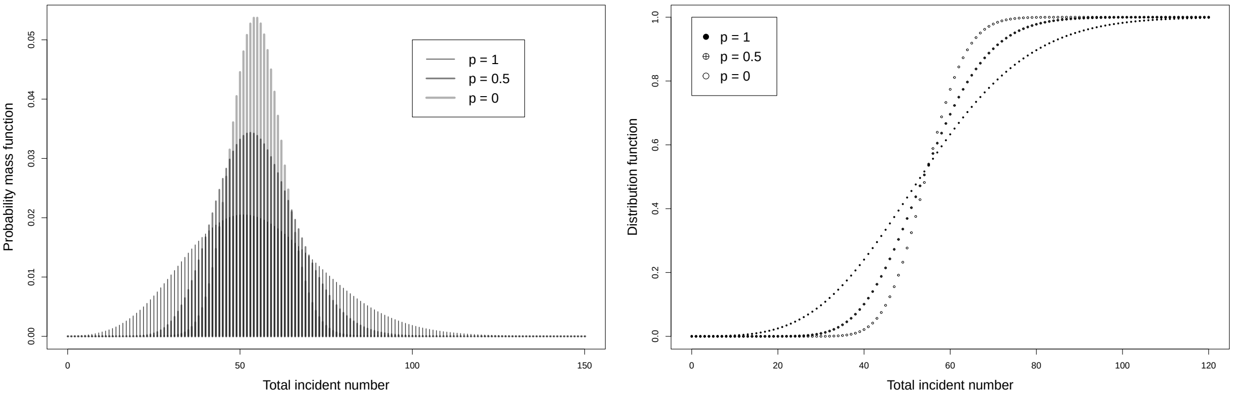

(a) Probability mass function of $S(T)$ for $p \in \{0, 0.5, 1\}$.

(b) Cumulative distribution function of $S(T)$ for $p \in \{0, 0.5, 1\}$.

Figure 3: Panel 3a shows the p.m.f. of the total incident number for rates $\boldsymbol{\lambda}$ as in Table 1, $T = 1$, and resulting rates $\widetilde{\boldsymbol{\lambda}}$ for $p \in \{0, 0.5\}$. Figure 3b analogously plots the c.d.f.s, illustrating that while the c.d.f.s are not ordered in the sense $F_{S_{\widetilde{M}}(T)}(x) \geq F_{S_M(T)}(x)$, $\forall x \geq 0$, there is a threshold value x_0 s.t. this ordering holds (exactly) for large values $x > x_0 \geq 0$, i.e. the so-called *single-crossing condition* is fulfilled here. In the actuarial context, one is typically interested in high quantiles of the loss distribution ($\mathbf{VaR}_{1-\gamma}$ for $(1-\gamma)$ close to 1), i.e. the region where in this case it holds for the quantile functions $F^{\leftarrow}_{S_{\widetilde{M}}(T)}(1-\gamma) \leq F^{\leftarrow}_{S_M(T)}(1-\gamma)$, leading to the observations for the portfolio risk measure discussed in this section.

Definition 2 (Bivariate Cuadras–Augé copula, [22], p. 9)**.** *For* $\alpha \in [0, 1]$*, let* $C_\alpha : [0, 1]^2 \mapsto [0, 1]$ *be defined by*

$$C_\alpha(u_1, u_2) := \min\{u_1, u_2\} \max\{u_1, u_2\}^{1-\alpha}, \quad u_1, u_2 \in [0, 1].$$

Remark 2 (Tail dependence coefficients of Cuadras–Augé (survival) copula ([22], p. 34f))**.** *For a bivariate Cuadras–Augé copula* C_α*, the tail dependence coefficients are given by*

$$UTD_{C_\alpha} = \alpha, \quad LTD_{C_\alpha} = \mathbf{1}_{\{\alpha=1\}}.$$

Note that in general for a copula C *and its survival copula* $\hat{C}$*, it holds (provided existence) that* $UTD_C = LTD_{\hat{C}}$ *and* $LTD_C = UTD_{\hat{C}}$*, respectively.*

This means for the comparison of the instantaneous joint loss arrival rate in (6), we are interested in comparing the parameter α (as in (7)) for models (M) and $(\widetilde{M})$.

Remark 3 ($LTD_{\hat{C}_\alpha}$ for constant $\boldsymbol{\lambda}$)**.** *Assume* $\lambda^{|I|=i} \equiv \bar{\lambda} > 0$, $\forall i \in \{1, \ldots, K\}$*. Then, in model* (M) *the lower tail dependence coefficient of the bivariate copula of* (τ_i, τ_j) *is given by*

$$LTD_{\hat{C}_\alpha} = \alpha = \frac{2}{3},$$

and the instantaneous joint loss arrival rate in (6) *is given by*

$$\lim_{T \searrow 0} \frac{\mathbb{P}(\tau_i \leq T, \tau_j \leq T)}{T} = \lambda^i \alpha = \frac{\bar{\lambda}(K+1)}{2} \cdot \frac{2}{3} = \frac{\bar{\lambda}(K+1)}{3}.$$

Lemma 2 (Relation of $LTD_{\hat{C}_\alpha}$ for models (M) and $(\widetilde{M})$)**.** *Let (M) be an exchangeable model as in Assumption 1 with any vector of arrival rates $\boldsymbol{\lambda}$ and let $(\widetilde{M})$ be the corresponding model according to Definition 1. Let α and $\widetilde{\alpha}$ be the respective parameters of the bivariate survival copulas of (any two) first-arrival times (τ_i, τ_j) as given in (7). Then, it holds that $\widetilde{\alpha} \leq \alpha$ and more specifically, under Assumption 2,*

$$\widetilde{\alpha} = p^2 \alpha$$

for any $p \in [0,1]$.

Lemma 2 implies that in model $(\widetilde{M})$, by omitting information about common events according to Assumption 2, the instantaneous joint loss arrival rate for any two companies in the portfolio is underestimated by a factor of p^2, which intuitively makes sense, as this factor indicates the probability of independently not overlooking a joint event in two companies.

Stochastic ordering and coherent risk measures

Above, we have observed exemplarily that the portfolio risk when measured by Value-at-Risk (at 'relevant' levels in an actuarial context, see the remark about the single-crossing condition above and illustration in Figure 3b) is underestimated in a model with missing information $(\widetilde{M})$ compared to an original model (M). Another important risk measure is *Expected Shortfall* (at level $(1-\gamma)$), in the following denoted $\mathbf{ES}_{1-\gamma}(X)$ for a r.v. X in the actuarial context, defined as (see e.g. [5]):

$$\mathbf{ES}_{1-\gamma}(X) = \frac{1}{\gamma} \int_{1-\gamma}^{1} \mathbf{VaR}_z(X) \mathrm{d}z, \tag{8}$$

where $\mathbf{VaR}_z(X)$ is defined in (5). It is well-known that $\mathbf{ES}_{1-\gamma}$ possesses in a certain sense preferable analytical properties compared to $\mathbf{VaR}_{1-\gamma}$, in particular $\mathbf{ES}_{1-\gamma}$ is a *coherent* risk measure.[16][17] The fact of $\mathbf{ES}_{1-\gamma}$ being coherent allows to draw some interesting theoretical conclusions for the present study presented below in Corollary 1. As a basis, we use the more general observation on the stochastic ordering of compound Poisson random variables summarized in the following theorem.

Theorem 1 (Increasing convex order for specific compound Poisson distributions)**.** *Let $L > 0$ and $\ell \in \mathbb{N}$ and consider two independent homogeneous Poisson processes with intensities*

[16] See the seminal work of [8] for the definition and properties of coherent risk measures and e.g. [16] for a collection of proofs of the coherence of expected shortfall.

[17] Note that the term 'expected shortfall' is often simply used interchangeably with 'average / tail / conditional Value-at-Risk' or 'tail conditional expectation', which are in turn usually used synonymously. In an actuarial context, the most well-known definition is $\mathbf{TVaR}_{1-\gamma}(X) = \mathbb{E}[X | X \geq \mathbf{VaR}_{1-\gamma}(X)]$, i.e. the expected loss given that a loss at least equal to the Value-at-Risk occurs. However, many equivalencies between the above risk measures, and in particular the coherence of the risk measures other than $\mathbf{ES}_{1-\gamma}$ as defined in (8), only hold if X follows a continuous distribution; see [5] for a detailed discussion. As in the context of this work, discrete underlying distributions (of incident numbers) occur, we therefore only consider $\mathbf{ES}_{1-\gamma}$.

$\lambda > 0$ and $\widetilde{\lambda} := \ell\,\lambda > 0$, denoted $N(t) := (N(t))_{t\geq 0}$ and $\widetilde{N}(t)$, respectively. For any fixed $T > 0$, let

$$S(T) := \sum_{i=1}^{N(T)} L = L\,N(T) \qquad \textit{and} \qquad \widetilde{S}(T) = \sum_{i=1}^{\widetilde{N}(t)} \frac{L}{\ell} = \frac{L}{\ell}\widetilde{N}(T).$$

Then, $\mathbb{E}\big[S(T)\big] = \mathbb{E}\big[\widetilde{S}(T)\big]$ and

$$S(T) \geq_{icx} \widetilde{S}(T), \tag{9}$$

where $\geq_{icx}$ denotes 'increasing convex order'.

Proof. See Appendix A.[18] □

Remark 4 (Notes to Theorem 1)**.**

1. *Note that $S(T) \geq_{icx} \widetilde{S}(T)$ and $\mathbb{E}\big[S(T)\big] = \mathbb{E}\big[\widetilde{S}(T)\big]$ is equivalent to $S(T) \geq_{cx} \widetilde{S}(T)$ (*'convex order'*), see [25], Theorem 1.5.3.*

2. *In actuarial science, a perhaps more common, synonymous name for* 'increasing convex order' *($\geq_{icx}$) is* 'stop-loss order' *($\geq_{sl}$), which stems from an important characterization of $\geq_{icx}$ by the so-called stop-loss transforms (see [25], Theorem 1.5.7):*

$$X \leq_{icx} Y \iff \mathbb{E}\big[(X-t)_+\big] \leq \mathbb{E}\big[(Y-t)_+\big] \ \ \forall t \in \mathbb{R}. \tag{10}$$

3. *Note that $S(T)$ and $\widetilde{S}(T)$ can be interpreted as two collective risk models with equal expected total claims amount $\mathbb{E}\big[S(T)\big] = \mathbb{E}\big[\widetilde{S}(T)\big]$, where*

   - ⋄ *$S(T)$ is the total claims amount from a model with relatively few, large losses (of deterministic size $L > 0$), and*
   - ⋄ *$\widetilde{S}(T)$ is the total claims amount from a model with relatively many, small losses (of deterministic size $0 < \frac{L}{\ell} < L$).*

   Thus, Theorem 1 states that the model with on average many (independent) small losses is preferable ('less risky'*) in the sense of increasing convex order compared to a model with equal expected claims amount and on average few (independent) large losses.*

Corollary 1 (Expected Shortfall for models (M) and $(\widetilde{M})$)**.** *Let $\mathbf{ES}_{1-\gamma}(\cdot)$ denote Expected Shortfall as in (8) and let $S_M(T)$ and $S_{\widetilde{M}}(T)$ denote the total incident number in the portfolio under models (M) and $(\widetilde{M})$, respectively, until a fixed time $T > 0$. Then, for any $T > 0$ and any $\gamma \in (0,1)$, it holds*

$$\mathbf{ES}_{1-\gamma}\big(S_M(T)\big) \geq \mathbf{ES}_{1-\gamma}\big(S_{\widetilde{M}}(T)\big). \tag{11}$$

Proof. See Appendix A. □

This implies that by omitting information about common events, the portfolio risk is necessarily underestimated when using expected shortfall (or any other coherent risk measure).

[18] Somewhat surprising to us, we did not find the (or a correspondent) statement of the theorem in the literature, hence, for completeness we provide an elementary proof in the Appendix.

4 Conclusion

When insurers started to develop actuarial models for cyber risk, they soon emphasized that one major challenge is the lack of adequate data to calibrate and backtest their models. Many classical actuarial models are based on the assumption of independence between losses and historical data is mainly used to draw inference about individual policyholders' loss distributions (i.e. the parameters of their loss frequency and severity distribution for a certain risk). Indeed, this is sufficient in markets where the claims are independent. Risk assessment and claims settlement therefore usually take into account this individual client-specific information. However, in the case of cyber, collecting such individual information alone is not sufficient, as not only parameters of the individual (marginal) loss distributions, but also those of an adequate model of dependence, have to be calibrated. This is only possible if information about dependence between historical claims, i.e. that losses may have stemmed from the same cause, is systematically collected.

This article has used a stylized mathematical model to highlight the effects on portfolio risk measurement if information on common events is fully or partly discarded. In practice, and we have to raise a big warning sign here, the resulting underestimation of accumulation risk would only become evident too late, namely once a (to-be-avoided) extreme portfolio loss has occurred.

The urgent practical implications for insurers are evident: As outlined in Section 2.1, actuarial modelling of cyber cannot be regarded as an isolated challenge, but as one interconnected step in the insurance value chain. Actuaries therefore must be in continuous exchange with other stakeholders, in particular legal experts (regarding insurability of cyber, product design, and requirements on the collection of claims settlement data) and information security experts. The central importance of the latter group for the actuarial modelling of cyber can hardly be overstated; their expertise is essential in tackling important challenges such as how to include an extensive qualitative assessment of a company's IT landscape, including existing security provisions, into a stochastic actuarial model.

Only continuous interdisciplinary cooperation will allow to develop a holistic approach which allows insurers to proactively steer their cyber underwriting activities without exposing themselves to potentially starkly underestimated levels of accumulation risk.

Acknowledgement

This work was conducted at ERGO Center of Excellence in Insurance at Technical University of Munich and the authors would like to thank ERGO Group AG for supporting this research and providing access to insurance industry expertise.

References

[1] https://privacyrights.org/data-breaches.

[2] https://www.advisenltd.com/data/cyber-loss-data/.

[3] *Hurricane Katrina 10.* https://www.agcs.allianz.com/news-and-insights/reports/lessons-learned-from-hurricane-katrina.html, 2015.

[4] *Tianjin Blast Could Be Largest Marine Insurance Loss Ever.* https://maritime-executive.com/article/tianjin-blast-could-be-largest-marine-insurance-loss-ever, 2016.

[5] C. Acerbi and D. Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, 2002.

[6] Advisen. *2018 Cyber Guide: The Ultimate Guide to Cyber Service Providers.* https://www.advisenltd.com/media/reports/cyber-guide/.

[7] Allianz Global Corporate & Specialty. *A Guide to Cyber Risk: Managing the Impact of Increasing Interconnectivity.* https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/AGCS-Cyberrisk-report.pdf, 2015.

[8] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

[9] K. Awiszus, T. Knispel, I. Penner, G. Svindland, A. Voß, and S. Weber. Modeling and pricing cyber insurance: Idiosyncratic, systematic, and systemic risks. *European Actuarial Journal*, 2023.

[10] Baldwin, A., Gheyas, I., Ioannidis, C., Pym, D., and Williams, J. Contagion in cyber security attacks. *Journal of the Operational Research Society*, 68(7):780–791, 2017.

[11] Bandyopadhyay, T., Mookerjee, V., and Rao, R. Why IT managers don't go for cyber-insurance products. *Communications of the ACM*, 52(11):68, 2009.

[12] Böhme, R., Laube, S., and Riek, M. A fundamental approach to cyber risk analysis. *Variance*, 11(2), 2018.

[13] Daley, D.J. and Vere-Jones, D. *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods.* Springer, 2. ed. edition, 2003.

[14] Edwards, B., Hofmeyr, S., and Forrest, S. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity*, 2(1):3–14, 2016.

[15] Eling, M. and Wirfs, J.H. What are the actual costs of cyber risk events? *European Journal of Operational Research*, 272(3):1109–1119, 2019.

[16] P. Embrechts and R. Wang. Seven proofs for the subadditivity of expected shortfall. *Dependence Modeling*, 3(1), 2015.

[17] Fahrenwaldt, M., Weber, S., and Weske, K. Pricing of cyber insurance contracts in a network model. *ASTIN Bulletin*, 48(3):1175–1218, 2018.

[18] Murchu L. Falliere, N. and E. Chien. W32.stuxnet dossier. *Symantec*, 2010.

[19] Gesamtverband der Deutschen Versicherungswirtschaft e.V. *Unverbindlicher Muster-Fragebogen zur Risikoerfassung im Rahmen von Cyber-Versicherungen für kleine und mittelständische Unternehmen. (Unverbindliche Bekanntgabe des Gesamtverbandes der Deutschen Versicherungswirtschaft e.V. (GDV) zur fakultativen Verwendung. Abweichende Vereinbarungen sind möglich.)*, 2019.

[20] Herath, H. and Herath, T. Copula-based actuarial model for pricing cyber-insurance policies. *Insurance Markets and Companies*, 2(1), 2011.

[21] J. Kingman. *Poisson processes*, volume 3 of *Oxford studies in probability*. Clarendon Press, Oxford, 1993.

[22] Mai, J. and Scherer, M. *Simulating copulas: Stochastic models, sampling algorithms, and applications*, volume vol. 6 of *Series in quantitative finance*. World Scientific Publishing, New Jersey and London and Singapore, 2nd edition edition, 2017.

[23] Marotta, A., Martinelli, F., Nanni, S., Orlando, A., and Yautsiukhin, A. Cyber-insurance survey. *Computer Science Review*, 24:35–61, 2017.

[24] Mukhopadhyay, A., Chatterjee, S., Saha, D., Mahanti, A., and Sadhukhan, S. Cyber-risk decision models: To insure it or not? *Decision Support Systems*, 56:11–26, 2013.

[25] A. Müller and D. Stoyan. *Comparison Methods for Stochastic Models and Risks*, volume 389. Wiley, 2002.

[26] Munich Re. Cyber insurance: Risks and trends 2020, 2020.

[27] H. Panjer. Recursive evaluation of a family of compound distributions. *ASTIN Bulletin: The Journal of the IAA*, 12(1):22–26, 1981.

[28] Peng, C., Xu, M., Xu, S., and Hu, T. Modeling and predicting extreme cyber attack rates via marked point processes. *Journal of Applied Statistics*, 44(14):2534–2563, 2017.

[29] Peng, C., Xu, M., Xu, S., and Hu, T. Modeling multivariate cybersecurity risks. *Journal of Applied Statistics*, 45(15):2718–2740, 2018.

[30] S. Romanosky. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity*, 2(2):121–135, 2016.

[31] Romanosky, S., Ablon, L., Kuehn, A., and Jones, T. Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity*, 5(1):117, 2019.

[32] B. Schneier. The story behind the Stuxnet virus. *Forbes*, 07.10.2010.

[33] tenable. tenable's 2021 threat landscape retrospective. *tenable Research*, 2021.

[34] W. Turton and J. Robertson. Microsoft attack blamed on china morphs into global crisis. *Bloomberg*, 07.03.2021.

[35] Xu, M., Da, G., and Xu, S. Cyber epidemic models with dependences. *Internet Mathematics*, 11(1):62–92, 2015.

[36] G. Zeller and M. Scherer. A comprehensive model for cyber risk based on marked point processes and its application to insurance. *European Actuarial Journal*, 12(1):33–85, 2022.

A Appendix A

Proof of Lemma 1

Proof of Lemma 1. Starting from Definition 1, we observe that the new marginal rates for any $\ell \in \{1, \dots, K\}$ are given by

$$\begin{aligned}
\widetilde{\lambda}^{\ell} = \sum_{i=1}^{K} \frac{i}{K} \widetilde{\lambda}^{|I|=i} &= \frac{1}{K}\lambda^{|I|=1} + \frac{1}{K}\Bigg[\sum_{i=2}^{K} \lambda^{|I|=i}\Big[i\big(f_{\text{Bin}}(0;i,p) + f_{\text{Bin}}(1;i,p)\big) \\
&\quad + \sum_{j=2}^{\max(i-1,2)} (i-j) f_{\text{Bin}}(j;i,p)\Big]\Bigg] + \sum_{i=2}^{K} \frac{i}{K} \sum_{j=i}^{K} \lambda^{|I|=j} f_{\text{Bin}}(i;j,p) \\
&= \frac{1}{K}\lambda^{|I|=1} + \frac{1}{K}\Big[\underbrace{\sum_{i=2}^{K} \lambda^{|I|=i} i\big(f_{\text{Bin}}(0;i,p) + f_{\text{Bin}}(1;i,p)\big)}_{(S1)} \\
&\quad + \underbrace{\sum_{i=2}^{K} \lambda^{|I|=i} \sum_{j=2}^{\max(i-1,2)} (i-j) f_{\text{Bin}}(j;i,p)}_{(S2)} + \underbrace{\sum_{i=2}^{K} i \sum_{j=i}^{K} \lambda^{|I|=j} f_{\text{Bin}}(i;j,p)}_{(S3)}\Big].
\end{aligned}$$

It remains to show that the sum in the square bracket equals $\sum_{j=2}^{K} j\lambda^{|I|=j}$. Reversing the order of summation in $(S3)$ and renaming $i \leftrightarrow j$ in the remaining terms yields

$$\begin{aligned}
&\Big[(S1) + (S2) + (S3)\Big] = \sum_{j=2}^{K} \lambda^{|I|=j} j\big(f_{\text{Bin}}(0;j,p) + f_{\text{Bin}}(1;j,p)\big) \\
&+ \sum_{j=2}^{K} \lambda^{|I|=j} \sum_{i=2}^{\max(j-1,2)} (j-i) f_{\text{Bin}}(i;j,p) + \sum_{j=2}^{K} \lambda^{|I|=j} \sum_{i=2}^{j} i f_{\text{Bin}}(i;j,p) \\
&= \sum_{j=2}^{K} \lambda^{|I|=j} j\big(f_{\text{Bin}}(0;j,p) + f_{\text{Bin}}(1;j,p)\big) + \sum_{j=2}^{K} \lambda^{|I|=j} \Big(\sum_{i=2}^{j-1} j f_{\text{Bin}}(i;j,p) + j f_{\text{Bin}}(j;j,p)\Big) \\
&= \sum_{j=2}^{K} j\lambda^{|I|=j} \underbrace{\sum_{i=0}^{j} f_{\text{Bin}}(i;j,p)}_{=1} = \sum_{j=2}^{K} j\lambda^{|I|=j}.
\end{aligned}$$

□

Proof of Remark 3

Proof of Remark 3. Note that due to the properties of the Binomial coefficient, it holds that

$$\begin{aligned}
\frac{\binom{K-1}{i}}{\binom{K}{i+1}} &= \frac{i+1}{K}, \\
\frac{\binom{K-2}{i}}{\binom{K}{i+1}} &= \frac{\frac{K-1-i}{K-1}\binom{K-1}{i}}{\binom{K}{i+1}} = \frac{K-1-i}{K-1}\cdot\frac{i+1}{K} = \frac{(K-(i+1))(i+1)}{K(K-1)}.
\end{aligned}$$

Inserting this into the expression in (7) yields

$$\begin{aligned}
\alpha = 1 - \frac{\sum_{i=0}^{K-2}\binom{K-2}{i}\frac{1}{\binom{K}{i+1}}\lambda^{|I|=i+1}}{\sum_{i=0}^{K-1}\binom{K-1}{i}\frac{1}{\binom{K}{i+1}}\lambda^{|I|=i+1}} &= 1 - \frac{\frac{1}{K(K-1)}\bar{\lambda}\sum_{i=0}^{K-2}(K-(i+1))(i+1)}{\frac{1}{K}\bar{\lambda}\sum_{i=0}^{K-1}(i+1)} \\
= 1 - \frac{1}{K-1}\frac{\sum_{i=1}^{K-1}(K-i)i}{\sum_{i=1}^{K} i} = 1 - \frac{1}{K-1}\frac{\frac{1}{6}K(K+1)(K-1)}{\frac{1}{2}K(K+1)} &= 1 - \frac{2}{6} = \frac{2}{3}.
\end{aligned}$$

For the marginal rates λ^i in (2), it holds

$$\lambda^i = \sum_{k=1}^{K}\frac{k}{K}\lambda^{|I|=k} = \frac{\bar{\lambda}}{K}\sum_{k=1}^{K} k = \frac{\bar{\lambda}}{K}\frac{K(K+1)}{2} = \frac{\bar{\lambda}(K+1)}{2},$$

implying the remark. □

Proof of Lemma 2

Proof of Lemma 2. By definition, α and $\widetilde{\alpha}$ are given by

$$\alpha = 1 - \frac{\sum_{i=1}^{K-1}\binom{K-2}{i-1}\lambda_i}{\sum_{i=1}^{K}\binom{K-1}{i-1}\lambda_i} =: 1 - \frac{Z_\alpha}{N_\alpha}, \quad \widetilde{\alpha} = 1 - \frac{\sum_{i=1}^{K-1}\binom{K-2}{i-1}\tilde{\lambda}_i}{\sum_{i=1}^{K}\binom{K-1}{i-1}\tilde{\lambda}_i} =: 1 - \frac{Z_{\widetilde{\alpha}}}{N_{\widetilde{\alpha}}},$$

where $\lambda_i = \frac{\lambda^{|I|=i}}{\binom{K}{i}}$ and $\widetilde{\lambda}_i = \frac{\widetilde{\lambda}^{|I|=i}}{\binom{K}{i}}$.

We use the following properties of the Binomial coefficient and the Binomial distribution

$$\binom{K-1}{i-1} = \frac{i}{K}\binom{K}{i}, \tag{BIN1}$$

$$\binom{K-2}{i-1} = \frac{K-i}{K-1}\binom{K-1}{i-1} = \frac{K-i}{K-1}\frac{i}{K}\binom{K}{i}, \tag{BIN2}$$

$$\binom{K-2}{i-2} = \binom{K-1}{i-1} - \binom{K-2}{i-1}. \tag{BIN3}$$

$$X \sim \text{Binom}(K,p) \implies \mathbb{E}[X] = Kp, \tag{BIN4}$$

$$X \sim \text{Binom}(K,p) \implies \mathbb{E}\big[X^2\big] = Kp(1-p) + K^2p^2. \tag{BIN5}$$

This implies the following auxiliary result:

$$N_\alpha - Z_\alpha = \sum_{i=1}^{K} \binom{K-1}{i-1}\lambda_i - \sum_{i=1}^{K-1} \binom{K-2}{i-1}\lambda_i \overset{\text{(BIN3)}}{=} \lambda_K + \sum_{i=2}^{K-1} \binom{K-2}{i-2}\lambda_i = \sum_{i=2}^{K} \binom{K-2}{i-2}\lambda_i. \tag{12}$$

Furthermore, it holds that $N_\alpha = N_{\widetilde{\alpha}}$, as

$$N_{\widetilde{\alpha}} = \sum_{i=1}^{K} \binom{K-1}{i-1}\widetilde{\lambda}_i = \sum_{i=1}^{K} \binom{K-1}{i-1}\frac{\widetilde{\lambda}^{|I|=i}}{\binom{K}{i}} \overset{\text{(BIN1)}}{=} \sum_{i=1}^{K} \frac{i}{K}\widetilde{\lambda}^{|I|=i} \overset{\text{Lemma 1}}{=} \sum_{i=1}^{K} \frac{i}{K}\lambda^{|I|=i} = N_\alpha. \tag{13}$$

We will show that for $Z_{\widetilde{\alpha}}$ it holds that

$$Z_{\widetilde{\alpha}} = \lambda_1 + \sum_{i=2}^{K} \lambda_i \Big[\binom{K-1}{i-1} - \binom{K-2}{i-2}p^2\Big]. \tag{$*$}$$

This implies the claim, as one can rewrite

$$\begin{aligned} Z_{\widetilde{\alpha}} &= \lambda_1 + \sum_{i=2}^{K} \lambda_i \Big[\binom{K-1}{i-1} - \binom{K-2}{i-2}p^2\Big] = \sum_{i=1}^{K} \binom{K-1}{i-1}\lambda_i - p^2 \sum_{i=2}^{K} \binom{K-2}{i-2}\lambda_i \\ &\overset{(12),(13)}{=} N_\alpha - p^2(N_\alpha - Z_\alpha). \end{aligned} \tag{14}$$

From this it follows

$$\widetilde{\alpha} = 1 - \frac{Z_{\widetilde{\alpha}}}{N_{\widetilde{\alpha}}} \overset{(13),(14)}{=} 1 - \frac{N_\alpha - p^2(N_\alpha - Z_\alpha)}{N_\alpha} = 1 - \Big[1 - p^2\Big(1 - \frac{Z_\alpha}{N_\alpha}\Big)\Big] = p^2\alpha.$$

To show $(*)$, we rewrite (3) as

$$\begin{aligned} \widetilde{\lambda}^{|I|=1} &= \lambda^{|I|=1} + \sum_{i=2}^{K} \lambda^{|I|=i}\Big[i\big(f_{\text{Bin}}(0;i,p) + f_{\text{Bin}}(1;i,p)\big) + \sum_{j=2}^{\max(i-1,2)} (i-j) f_{\text{Bin}}(j;i,p)\Big] \\ &= \lambda^{|I|=1} + \lambda^{|I|=2} 2\underbrace{\big(f_{\text{Bin}}(0;2,p) + f_{\text{Bin}}(1;2,p)\big)}_{(1-p^2)} + \sum_{i=3}^{K} \lambda^{|I|=i}\Big[\underbrace{i\sum_{j=0}^{i-1} f_{\text{Bin}}(j;i,p)}_{i(1-p^i)} - \underbrace{\sum_{j=2}^{i-1} j f_{\text{Bin}}(j;i,p)}_{\pm \sum_{j=0,1,i} j f_{\text{Bin}}(j;i,p)}\Big] \\ &\overset{\text{(BIN4)}}{=} \lambda^{|I|=1} + \lambda^{|I|=2} 2(1-p^2) + \sum_{i=3}^{K} \lambda^{|I|=i}\big[i - ip^i - \big(ip - ip(1-p)^{i-1} - ip^i\big)\big] \\ &= \lambda^{|I|=1} + \lambda^{|I|=2} 2(1-p^2) + \sum_{i=3}^{K} \lambda^{|I|=i} i(1 - p + p(1-p)^{i-1}) \\ &= \lambda^{|I|=1} + \sum_{i=2}^{K} \lambda^{|I|=i} i(1 - p + p(1-p)^{i-1}). \end{aligned}$$

Changing to the rates $\widetilde{\lambda}_1 = \frac{\widetilde{\lambda}^{|I|=1}}{K}$ (LHS) and $\lambda_i = \frac{\lambda^{|I|=i}}{\binom{K}{i}}$ (RHS) yields

$$\widetilde{\lambda}_1 = \lambda_1 + \sum_{i=2}^{K} \frac{i}{K}\binom{K}{i}\lambda_i(1-p+p(1-p)^{i-1}) \overset{(BIN1)}{=} \lambda_1 + \sum_{i=2}^{K}\binom{K-1}{i-1}\lambda_i(1-p+p(1-p)^{i-1}),$$

i.e. for fixed $i \in \{2,\dots,K\}$, the coefficient of λ_i from $\widetilde{\lambda}_1$, which appears in $Z_{\widetilde{\alpha}}$ with factor $\binom{K-2}{0} = 1$ is given by $\binom{K-1}{i-1}(1-p+p(1-p)^{i-1})$. Analogously, the coefficients of λ_i from $\sum_{j=2}^{K-1}\widetilde{\lambda}_j$, scaled by $\binom{K-2}{j-1}$, are illustrated as the column sums in Table 3 and given by

$$\begin{aligned}
&\underbrace{\lambda_i\binom{K}{i}}_{\lambda_i\binom{K}{i}=\lambda^{|I|=i}} \sum_{j=2}^{i}\binom{i}{j}p^j(1-p)^{i-j}\underbrace{\binom{K-2}{j-1}}_{\text{Def. of } Z_\alpha}\underbrace{\frac{1}{\binom{K}{j}}}_{\lambda_j=\frac{\lambda^{|I|=j}}{\binom{K}{j}}}\\
&\overset{\text{(BIN2)}}{=} \lambda_i\binom{K}{i}\sum_{j=2}^{i}\frac{K-j}{K-1}\frac{j}{K}f_{\text{Bin}}(j;i,p) = \frac{\lambda_i\binom{K}{i}}{(K-1)K}\Big[K\sum_{j=2}^{i} j f_{\text{Bin}}(j;i,p) - \sum_{j=2}^{i} j^2 f_{\text{Bin}}(j;i,p)\Big]\\
&\overset{\text{(BIN4),(BIN5)}}{=} \frac{\lambda_i\binom{K}{i}}{(K-1)K}\Big[K(ip-ip(1-p)^{i-1}) - \big[ip(1-p)+i^2p^2-ip(1-p)^{i-1}\big]\Big]\\
&= \frac{\lambda_i\binom{K}{i}}{(K-1)K}\Big[Kip - Kip(1-p)^{i-1} - ip + ip^2 - i^2p^2 + ip(1-p)^{i-1}\Big]\\
&= \frac{\lambda_i\binom{K}{i}}{(K-1)K}\Big[(K-1)ip - (K-1)ip(1-p)^{i-1} - (i-1)ip^2\Big]\\
&= \lambda_i\Big[\binom{K}{i}\frac{i}{K}p - \binom{K}{i}\frac{i}{K}p(1-p)^{i-1} - \binom{K}{i}\frac{i(i-1)}{K(K-1)}p^2\Big]\\
&\overset{\text{(BIN1)}}{=} \lambda_i\Big[\binom{K-1}{i-1}(p-p(1-p)^{i-1}) - \binom{K-2}{i-2}p^2\Big].
\end{aligned}$$

Thus, adding the coefficients of λ_i from $\widetilde{\lambda}_1$ and $\sum_{j=2}^{K}\widetilde{\lambda}_j\binom{K-2}{j-1}$ for each fixed $i \in \{2,\dots,K-1\}$ yields

$$\binom{K-1}{i-1}(1-p+p(1-p)^{i-1}) + \binom{K-1}{i-1}(p-p(1-p)^{i-1}) - \binom{K-2}{i-2}p^2 = \binom{K-1}{i-1} - \binom{K-2}{i-2}p^2,$$

which implies $(*)$ and therefore the claim. □

Proof of Theorem 1

Proof of Theorem 1.
Step 1: Increasing convex order for some discrete random variables
For an integer $K > 0$, consider a Bernoulli r.v. $Z \sim \text{Ber}(p)$, $p \in [0,1]$ and K i.i.d. copies of it denoted Z_i, $i \in \{1,\dots,K\}$.
Furthermore, consider the r.v.s X and Y defined as follows:

$$\begin{aligned}
X &= K\,Z,\\
Y &= \sum_{i=1}^{K} k_i\,Z_i, \quad i \in \{1,\dots,K\},
\end{aligned} \tag{15}$$

where $\boldsymbol{k} := (k_i)_{i\in\{1,\dots,K\}}$ is an $\mathbb{N}_0^K$-vector s.t. $\forall i:\ k_i \in \{0,\dots,K\}$ with $\sum_{i=1}^K k_i = \|\boldsymbol{k}\|_1 = K$. Assume w.l.o.g. $k_i \geq k_{i+1}$, $\forall i \in \{1,\dots,K-1\}$, and let $i^* := |\{k_i : k_i > 0\}|$, then the first i^* entries of $\boldsymbol{k}$ represent a *partition* of K (and the remaining entries equal 0).
It is obvious that for any r.v. Y as above

$$\mathbb{E}[Y] = \mathbb{E}[X] = Kp,$$

and we will now show that for any such Y it holds that

$$Y \leq_{icx} X$$

by using the following sufficient condition (the so-called *cut criterion* or *crossing condition*, see e.g. [25], p. 23): If for two r.v.s X and Y with c.d.f.s F_X and F_Y respectively, it holds that $\mathbb{E}[Y] \leq \mathbb{E}[X]$ and in addition, there exists $t_0 \in \mathbb{R}$ s.t.

$$\begin{aligned} F_Y(t) &\leq F_X(t) \quad \forall t < t_0,\\ F_Y(t) &\geq F_X(t) \quad \forall t \geq t_0, \end{aligned} \tag{16}$$

then this implies $Y \leq_{icx} X$.
Let us in the following exclude the trivial cases $p \in \{0,1\}$ and $\boldsymbol{k} = (K,0,\dots,0)$ as they lead to $F_X = F_Y$. Note that in all non-degenerate cases we have $i^* > 1$.
Then, for r.v.s X and Y as defined in (15), there exists $t_0 \in [1, K-1]$ s.t. the single-crossing condition is fulfilled:
For $t < 0$ and $t \geq K$, obviously $F_X(t) = F_Y(t)$.
For $t \in [0,1)$, we use that $p \in (0,1)$ and $i^* > 1$ to see

$$F_Y(t) = \mathbb{P}(Y = 0) = (1-p)^{i^*} < 1 - p = \mathbb{P}(X = 0) = F_X(t).$$

For $t \in (K-1, K)$, again with $p \in (0,1)$ and $i^* > 1$,

$$F_Y(t) = \mathbb{P}(Y \leq K-1) = 1 - \mathbb{P}(Y = K) = 1 - p^{i^*} > 1 - p = \mathbb{P}(X = 0) = F_X(t).$$

Lastly, note that

- $t \mapsto F_X(t)$ is constant for $t \in (0, K-1]$ at the level $F_X(t) \equiv 1 - p$.
- $F_Y(t)$ is monotone increasing (being a c.d.f.) for $t \in (0, K-1]$ with (non-negative) jumps at some of the $\{1,\dots,K-1\}$ and $F_Y(0+) = (1-p)^{i^*} < 1 - p < 1 - p^{i^*} = F_Y(K-1)$.

Thus, due to the monotonicity of F_Y, there must be a *unique* $t_0 \in [1, K-1]$ fulfilling (16).
Step 2: Implication for (compound) Poisson process setting
Now, fix a time horizon $T > 0$ and consider two independent homogeneous Poisson processes $N(t) := (N(t))_{t\geq 0}$ with rate $\lambda > 0$ and $\widetilde{N}(t) := (\widetilde{N}(t))_{t\geq 0}$ with rate $\ell\lambda > 0$, $\ell \in \mathbb{N}$. As $\widetilde{N}(t)$ can be understood (in the sense of being equal in distribution) as the superposition of ℓ independent Poisson processes $\widetilde{N}_j(t)$, $j \in \{1,\dots,\ell\}$, all of them with rate $\lambda > 0$ (see e.g. [21], p. 16), one can write $S(T)$ and $\widetilde{S}(T)$ as

$$\begin{aligned} S(T) &= \sum_{i=1}^{N(T)} L = L\ N(T),\\ \widetilde{S}(T) &= \sum_{i=1}^{\widetilde{N}(T)} \frac{L}{\ell} \overset{D}{=} \frac{L}{\ell} \sum_{j=1}^{\ell} \widetilde{N}_j(T). \end{aligned}$$

Due to the properties of the homogeneous Poisson process and by Wald's equation, it follows immediately that

$$\begin{aligned}
&N(T), \widetilde{N}_j(T) \sim \mathrm{Poi}(\lambda\, T),\ j \in \{1, \dots, \ell\},\\
&\mathbb{E}\big[N(T)\big] = \mathbb{E}\big[\widetilde{N}_j(T)\big] = \lambda\, T,\ j \in \{1, \dots, \ell\},\\
&\mathbb{E}\big[S(T)\big] = \mathbb{E}\big[\widetilde{S}(T)\big] = \lambda\, T\, L,
\end{aligned}$$

where $\mathrm{Poi}(\lambda)$ denotes the Poisson distribution with density $f_{\mathrm{Poi}(\lambda)}(k) = \frac{\lambda^k e^{-\lambda}}{k!},\ k \in \mathbb{N}_0,\ \lambda > 0$.
Now, consider the following random variables:

$$X^i = L\, \mathbf{1}_{\{N(T) \geq i\}} = \begin{cases} L & \text{if } N(T) \geq i, \\ 0 & \text{else,} \end{cases} \implies X^i = \begin{cases} L & \text{w.p. } 1 - \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j), \\ 0 & \text{w.p. } \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j), \end{cases}$$

$$Y^i_j = \frac{L}{\ell} \mathbf{1}_{\{\widetilde{N}_j(T) \geq i\}} = \begin{cases} \frac{L}{\ell} & \text{if } \widetilde{N}_j(T) \geq i, \\ 0 & \text{else,} \end{cases} \implies Y^i_j = \begin{cases} \frac{L}{\ell} & \text{w.p. } 1 - \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j), \\ 0 & \text{w.p. } \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j), \end{cases} \quad j \in \{1, \dots, \ell\}.$$

Note that X^i denotes the size of the i^{th} jump of the Poisson process $N(t)$ if it occurs until time T (of deterministic size $L > 0$ if the process jumps at least i times until time T, and of size 0 else), and analogously the ℓ independent random variables Y^i_j denote the sizes of the i^{th} jump of each of the independent Poisson processes $\widetilde{N}_j(t)$ if they occur until time T.
As the $Y^i_j,\ j \in \{1, \dots, \ell\}$, are independent, one can derive the density of their sum, denoted Y^i, from arguments borrowed from the Binomial law:

$$Y^i := \sum_{j=1}^{\ell} Y^i_j = \begin{cases} L & \text{w.p. } \Big(1 - \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big)^{\ell}, \\ \frac{\ell-1}{\ell} L & \text{w.p. } \binom{\ell}{\ell-1} \Big(1 - \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big)^{\ell-1} \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j), \\ \cdots \\ \frac{1}{\ell} L & \text{w.p. } \binom{\ell}{1} \Big(1 - \sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big) \Big(\sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big)^{\ell-1}, \\ 0 & \text{w.p. } \Big(\sum\limits_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big)^{\ell}. \end{cases}$$

Note that this illustrates the fundamental difference between the two considered cases (process $N(t)$ vs. superposition of ℓ processes $\widetilde{N}_j(t)$): In the notation of a collective risk model, if the claim occurrences are driven by the process $N(t)$ (corresponding to relatively few events) and claim sizes are relatively large (i.e. of size L), either a large total claims amount occurs or no claim at all occurs for each jump. On the contrary, if claim occurrences are driven by the independent processes $\widetilde{N}_j(t)$ or equivalently their superposition $\widetilde{N}(t)$ (relatively many events) and claim sizes are relatively small (i.e. of size $\frac{L}{\ell}$), for a large total claims amount of size L from all the first (second, third, ...) jumps to occur, all ℓ processes $\widetilde{N}_j(t)$ independently need to jump at least once (twice, three times, ...); equivalently, ℓ independent jumps need to occur before time T in the superimposed process $\widetilde{N}(t)$. Likewise, to obtain no claim at all from the i^{th} jumps, any of the processes $\widetilde{N}_j(t)$ independently must not jump more than $(i-1)$ times; or equivalently, the superimposed process may not jump more than $(i-1)\ell$ times until T. Therefore, the probability of both large (i.e. size L) and no (size 0) total claims amounts is reduced, and probability mass is shifted to the intermediate cases that some (but not all or none) of the independent processes observe at least i jumps. As

$$\mathbb{E}[X^i] = \mathbb{E}[Y^i] = L\Big(1 - \sum_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)\Big)$$

– note that the weights for Y^i are akin to the density of a Binomial distribution with $N = \ell,\ p = 1 - \sum_{j=0}^{i-1} f_{\mathrm{Poi}(\lambda T)}(j)$ – for any $i \in \mathbb{N}$ the discrete random variables X^i and Y^i are akin to X and Y from the first part of the proof, X^i being a Bernoulli r.v. with positive mass only on the largest admissible value L and Y^i following a discrete density supported on the set of values $\{0, \frac{L}{\ell}, \cdots, \frac{(\ell-1)L}{\ell}, L\}$ with equal expectation. It follows from the above derivations that $X^i \geq_{icx} Y^i,\ i \in \mathbb{N}$. As (increasing) convex order is preserved under summation (this follows immediately from the transitivity of $\leq_{icx}$), this implies the statement of the theorem as

$$S(T) = \sum_{i\in\mathbb{N}} X^i \geq_{icx} \sum_{i\in\mathbb{N}} Y^i = \widetilde{S}(T).$$

Note that it is straightforward to again extend the result to a case where not all deterministic jump sizes corresponding to the ℓ arrival processes $\widetilde{N}_\ell(t)$ are equally of size $\frac{L}{\ell}$, but instead one replaces them by a collection $\{L_i\}_{i\in\{1,\ldots,\ell\}}$, such that $L_i > 0, \forall i \in \{1, \ldots, \ell\}$, and $\sum L_i = L$. □

Proof of Corollary 1

Proof of Corollary 1. It is a well-known result that for any two integrable r.v. X and Y, convex order is equivalent to the ordering of expected shortfall at all levels q, i.e.

$$Y \leq_{cx} X \iff \mathbf{ES}_q(Y) \leq \mathbf{ES}_q(X),\ \ \forall q \in (0,1),$$

see e.g. [16] and the references therein. Therefore, the statement of the corollary is equivalent to showing $S_M(T) \geq_{cx} S_{\widetilde{M}}(T)$. As from Lemma 1 (and linearity) it follows that $\mathbb{E}[S_M(T)] = \mathbb{E}[S_{\widetilde{M}}(T)]$, it is sufficient to show $S_M(T) \geq_{icx} S_{\widetilde{M}}(T)$ (see first point of Remark 4).
This follows immediately from Theorem 1: Recall that in model (M), the arrival rates for events of size $k \in \{1, \ldots, K\}$ are given by $\boldsymbol{\lambda} := (\lambda^{|I|=1}, \ldots, \lambda^{|I|=K})$ and that all arrivals are independent (from arrivals of events of the same or any other size). The total incident

number until time T can therefore again be written as a sum of K independent compound Poisson r.v.s:

$$S_M(T) \overset{D}{=} \sum_{k=1}^{K} kN_k(T),$$

where $N_k(t) := (N_k(t))_{t\geq 0}, k \in \{1,\ldots,K\}$, are independent Poisson processes with rates $\lambda^{|I|=k}$. In turn, for any k, the process $N_k(t)$ can (artificially) be understood as the superposition of $(k+1)$ independent Poisson processes $N_{k,j}(t) := (N_{k,j}(t))_{t\geq 0},\ j \in \{0,\ldots,k\}$, with rates $\lambda^{|I|=k}\ f_{\text{Bin}}(j;k,p)$, where in model (M) each of the arrivals of each of these processes is associated with a jump of size k.

Then, the total incident number from events of size k until time $T > 0$, denoted $S_k(T)$, and events of all sizes, denoted $S(T)$, are given by the following compound Poisson r.v.s, respectively:

$$S_k(T) = \sum_{j=0}^{k} kN_{k,j}(T) \overset{D}{=} kN_k(T), \quad S(T) = \sum_{k=1}^{K} S_k(T) = \sum_{k=1}^{K}\sum_{j=0}^{k} kN_{k,j}(T),$$

where $\overset{D}{=}$ denotes equality in distribution.

In model (M), for any $k \in \{1,\ldots,K\}$ each Poisson arrival process $N_{j,k}(t)$ is associated with jumps of size k. In model $(\widetilde{M})$, each arrival process $N_{j,k}(t)$ is replaced by several independent processes with equal Poisson rate, but associated with smaller jump sizes (which sum up to k), as represented in Table 4. From Theorem 1, it follows immediately that the compound incident number in the second column (model (M)) of each row dominates in increasing convex order the compound incident number of the corresponding processes in the third column (model $(\widetilde{M})$). By summing over all rows (recall that $\leq_{icx}$ is preserved under summation), the same holds for the compound incident number from each process $N_k(t),\ k \in \{1,\ldots,K\}$, in model (M) as compared to the overall compound incident number from all the corresponding independent processes in model $(\widetilde{M})$. By summing over all $k \in \{1,\ldots,K\}$, it follows that $S_M(T) \geq_{icx} S_{\widetilde{M}}(T)$ for any fixed $T > 0$ and thus the statement of the corollary. □

	λ_2	λ_3	$\cdots$	λ_{K-2}	λ_{K-1}	λ_K	$\binom{K-2}{j-1}$
$\widetilde{\lambda}_K$						$\binom{K}{K}\lambda_K p^K \frac{\binom{K}{K}}{\binom{K}{K}}$	
$\widetilde{\lambda}_{K-1}$					$\binom{K-1}{K-1}\lambda_{K-1} p^{K-1} \frac{\binom{K}{K-1}}{\binom{K}{K-1}}$ +	$\binom{K}{K-1}\lambda_K p^{K-1}(1-p) \frac{\binom{K}{K}}{\binom{K}{K-1}}$	$\binom{K-2}{K-2}$
$\widetilde{\lambda}_{K-2}$				$\binom{K-2}{K-2}\lambda_{K-2} p^{K-2} \frac{\binom{K}{K-2}}{\binom{K}{K-2}}$ +	$\binom{K-1}{K-2}\lambda_{K-1} p^{K-2}(1-p) \frac{\binom{K}{K-1}}{\binom{K}{K-2}}$ +	$\binom{K}{K-2}\lambda_K p^{K-2}(1-p)^2 \frac{\binom{K}{K}}{\binom{K}{K-2}}$	$\binom{K-2}{K-3}$
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
$\widetilde{\lambda}_3$		$\binom{3}{3}\lambda_3 p^3 \frac{\binom{K}{3}}{\binom{K}{3}}$ +	$\cdots$ +	$\binom{K-2}{3}\lambda_{K-2} p^3(1-p)^{(K-2)-3} \frac{\binom{K}{K-2}}{\binom{K}{3}}$ +	$\binom{K-1}{3}\lambda_{K-1} p^3(1-p)^{(K-1)-3} \frac{\binom{K}{K-1}}{\binom{K}{3}}$ +	$\binom{K}{3}\lambda_K p^3(1-p)^{K-3} \frac{\binom{K}{K}}{\binom{K}{3}}$	$\binom{K-2}{2}$
$\widetilde{\lambda}_2$	$\binom{2}{2}\lambda_2 p^2 \frac{\binom{K}{2}}{\binom{K}{2}}$ +	$\binom{3}{2}\lambda_3 p^2(1-p) \frac{\binom{K}{3}}{\binom{K}{2}}$ +	$\cdots$ +	$\binom{K-2}{2}\lambda_{K-2} p^2(1-p)^{(K-2)-2} \frac{\binom{K}{K-2}}{\binom{K}{2}}$ +	$\binom{K-1}{2}\lambda_{K-1} p^2(1-p)^{(K-1)-2} \frac{\binom{K}{K-1}}{\binom{K}{2}}$ +	$\binom{K}{2}\lambda_K p^2(1-p)^{K-2} \frac{\binom{K}{K}}{\binom{K}{2}}$	$\binom{K-2}{1}$

Table 3: The table illustrates the calculation of the coefficient of each $\lambda_i,\ i \in \{2, \ldots, K-1\}$, in $Z_{\widetilde{\alpha}}$. In each row, $\widetilde{\lambda}_j$ is calculated based on the definition of model $(\widetilde{M})$ given in (4), where $\widetilde{\lambda}_j = \frac{\widetilde{\lambda}^{|I|=j}}{\binom{K}{j}}$ and $\lambda_i = \frac{\lambda^{|I|=i}}{\binom{K}{i}}$ are substituted (leading to the last fraction of Binomial coefficients in each entry). The coefficient of each $\lambda_i,\ i \in \{2, \ldots, K-1\}$, in $Z_{\widetilde{\alpha}}$ is given by the scalar product of the i^{th} column and the very last column which lists the factors $\binom{K-2}{j-1}$ from the definition of $Z_{\widetilde{\alpha}}$.

	Model (M): (Poisson rate, jump size)	Model $(\widetilde{M})$: (Poisson rate, jump size)	Interpretation
$N_{k,k}(T)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k;k,p),k\big)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k;k,p),k\big)$	All k joint arrivals recognized $\Longrightarrow$ process with jump size k is "replaced by" process with jump size k.
$N_{k,k-1}(T)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-1;k,p),k\big)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-1;k,p),k-1\big)$ $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-1;k,p),1\big)$	$k-1$ joint arrivals recognized $\Longrightarrow$ process with jump size k is replaced by two independent processes with jump sizes $k-1$ and 1, respectively.
$N_{k,k-2}(T)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-2;k,p),k\big)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-2;k,p),k-2\big)$ $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-2;k,p),1\big)$ $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(k-2;k,p),1\big)$	$k-2$ joint arrivals recognized $\Longrightarrow$ process with jump size k is replaced by three independent processes with jump sizes $k-2$, 1, and 1, respectively.
… …	… …	… …	… …
$N_{k,2}(T)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(2;k,p),k\big)$	$\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(2;k,p),k-2\big)$ $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(2;k,p),1\big)$ $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(2;k,p),1\big)$ … $\big(\lambda^{\lvert I\rvert=k}Tf_{\mathrm{Bin}}(2;k,p),1\big)$	2 joint arrivals recognized $\Longrightarrow$ process with jump size k is replaced by independent processes with jump sizes 2 (one process) and 1 ($k-2$ processes), respectively.
$N_{k,1}(T)$ + $N_{k,0}(T)$	$\big(\lambda^{\lvert I\rvert=k}T(f_{\mathrm{Bin}}(1;k,p)+f_{\mathrm{Bin}}(0;k,p)),k\big)$	$\big(\lambda^{\lvert I\rvert=k}T(f_{\mathrm{Bin}}(1;k,p)+f_{\mathrm{Bin}}(0;k,p)),1\big)$ … … $\big(\lambda^{\lvert I\rvert=k}T(f_{\mathrm{Bin}}(1;k,p)+f_{\mathrm{Bin}}(0;k,p)),1\big)$	1 or 0 joint arrivals recognized $\Longrightarrow$ process with jump size k is replaced by k independent processes, each with jump size 1.

Table 4: Comparison of the compound Poisson processes corresponding to models (M) and $(\widetilde{M})$ for any fixed $k \in \{1,\dots,K\}$: The arrival process $N_k(t)$ for events of size k, i.e. associated with jumps of size k, can be replaced by $\widetilde{(k+1)}$ independent processes with thinned rates according to the weights of a Binomial distribution. According to Definition 1, in model $(\widetilde{M})$, these processes are replaced by several independent processes, associated with smaller jump sizes adding up to k.

B.2 Cyber Insurance: An integral component of Cyber Risk Management [4]

Summary

This article is an invited contribution to the 'Yearbook 2021' of the Frankfurt Institute for Risk Management and Regulation (FIRM e.V.). The goal of FIRM e.V. is to *"promote teaching and research around the world of risk management and regulation"*[32] and to enable knowledge transfer between academia and industry. The yearbook represents a collection of spotlight expert contributions on current topics in risk management and regulation and is aimed at risk managers and other (C-level) decision makers in the area of risk and regulation in Germany.

The goal of this article is to foster an understanding of cyber insurance solutions as one building block of a holistic cyber risk management strategy. Therefore, the first part provides an overview of the German cyber insurance market at the time in terms of products and types of coverage (see, e.g., [124]) and emphasizes the potential to transcend mere risk transfer by offering services which help to prevent the manifestation of a cyber incident overall or mitigate its financial consequences. We also shed light on the factors which typically determine a company's risk rating and, hence, insurance premium, based on the extensive survey [121] and our own visualization of the German Insurance Association's (GDV) sample questionnaire for risk assessment for cyber insurance applicable to small and medium-sized enterprises ([78]).
The article furthermore describes the limits of insurability of cyber due to worrisome accumulation scenarios such as a long-lasting blackout (see, e.g., [106]) and the complications arising from the necessity to detect so-called *silent cyber* exposure in traditional lines of business.
The article concludes with an urgent appeal to abandon the beliefs that cyber is a merely technical topic isolated within IT teams and systems, or that risk avoidance is a feasible strategy for cyber as one's own exposure is negligible. Instead, two actionable conclusions are drawn: First, in today's interconnected world with business models heavily relying on functional internal and external IT infrastructures, cyber risk management should be regarded as a central part of enterprise risk management. Second, prospective cyber insurance buyers should not let themselves be deterred by the market's (perceived) lack of transparency as the risk assessment process for dedicated cyber coverage (independent of the final decision about purchasing insurance) provides a valuable opportunity for reflection of one's exposure to cyber threats and their potential (financial) consequences.

[32] See `https://www.firm.fm/en/`.

Reception

This article was awarded the *FIRM Yearbook Award 2021*.

Individual contributions

I am the main author of this article. I suggested the concrete topic and structure of the article based on an initial general discussion with my supervisor Matthias Scherer. I was responsible for the writing of the text and creation of the graphics. Matthias Scherer contributed feedback to improve both the content and presentation of the article.

Permission to include the article

Permission to include yearbook article in thesis

Title of article:

Die Cyberversicherung: ein integraler Bestandteil des Cyberrisikomanagements (DE)

Cyber Insurance: An integral component of Cyber Risk Management (ENG)

Authors:

Gabriela Zeller, Matthias Scherer

Published in:

FIRM e.V. Jahrbuch 2021, p. 22 - 26 and 124 - 128

Permission to include in thesis:

We grant the authors the right to include the above-mentioned article as printed in the yearbook in their own thesis. Furthermore, the authors may reproduce and make available their thesis, including the above-mentioned article, as required by their degree-awarding academic institution.

FIRM
Frankfurter Institut für
Risikomanagement und Regulierung

Frankfurt, 24.04.2023

Place, date

Signature, full name

Gesellschaft für Risikomanagement
und Regulierung e.V.
Schwarzwaldstraße 42
60528 Frankfurt am Main
Telefon +49 (0)69 87 40 20 00
Telefax +49 (0)69 87 40 20 09
info@firm.fm

Cyber Insurance: An integral component of Cyber Risk Management

Gabriela Zeller | Matthias Scherer

With increasing interconnectivity, dependence on information technology, and continuous expansion of IT systems, companies from all industry sectors have become aware of the exacerbating threat cyber risk poses to their business. Besides, many jurisdictions have started to impose fines for data protection violations. Likewise, the perception of cyber security has been moving from a merely technical topic, rooted within IT teams and systems, to being considered a larger business risk that needs to be incorporated into the ERM. In 2020, cyber incidents are ranked the number one peril to businesses worldwide, see [AGCS 2020]. For many companies, signing a dedicated cyber insurance policy constitutes one building block of a holistic risk management strategy. This article surveys today's cyber insurance market in terms of products and types of coverage and explains which factors determine a company's risk category and, hence, insurance premium. Furthermore, it describes limits of insurability of cyber risk and highlights the influence of `silent cyber' on traditional insurance policies.

Cyber Insurance Market: Status Quo

In 2015, the global market for cyber insurance was estimated to be worth around $2bn in premium, with US business accounting for approximately 90%. At the time, fewer than 10% of all companies had purchased cyber insurance, with typical buyers of cyber coverage coming from industries holding large volumes of personal data, e.g. healthcare and retail, or relying on digitalized technology processes, such as manufacturing and telecommunications. A rapid market growth was projected, with total premium reaching $20+ billion by 2025, see [AGCS 2015]. Seen from today, this estimate still seems realistic, with a global market size of around $7bn in 2020, see [Munich Re 2020a]. A testimony to the relative maturity of the US market is the yearly published `Advisen Cyber Guide,' which listed around 170 cyber service providers in the US in 2018. Among them were 43 carriers underwriting cyber risk, 23 brokers, and 15 organizations providing service in `Insurance Data and Analytics,' i.e. collecting data on cyber losses, conducting data-driven cyber analytics, and building probabilistic models for cyber losses, see [Advisen 2018]. Concerning the local German market, as of 2020 a comprehensive overview lists 14 insurers offering stand-alone cyber coverages, with cover limits ranging up until □25mn, see [Schonschek 2020].

Cyber Insurance Coverage: Combining risk transfer and risk mitigation

Companies considering purchasing cyber insurance are advised to first scrutinize their coverage requirements. Often, responses about most sought-after types of coverage are influenced by news of cyber losses experienced by others. Consequently, while a few years ago losses related to data breaches were the top concern for most companies, more recently fund transfer fraud / social engineering, cyber extortion / ransomware, and particularly cyber-related business interruption have moved to the fore, see [Advisen and PartnerRe 2018, Advisen and Zurich 2019].

The connection between cyber risk and business interruption (BI) is a particularly critical one, as the latter has led for many years the list of global top business risks only to be overtaken by the former in 2020. Furthermore, these risks are very much intertwined: Cyber incidents are the most feared BI-trigger and likewise, BI is the main cause of an economic loss after a cyber incident, see [AGCS 2019, 2020]. As businesses continue to rely on increasingly complex, global supply chain systems that in turn depend on the interconnectivity of systems and services, the potential (contingent) business interruption exposure is set to increase for businesses and insurers alike.

Once a company has identified its coverage needs, an adequate policy has to be selected from the offers on the market. Surveys indicate that this is often perceived as non-trivial; many prospective buyers report difficulties in understanding cyber insurance policies and criticize the lack of consistency between different carriers, see [Advisen and Zurich 2019].

Most standard policies cover financial losses originating from a compromise of one of the three classical information security protection goals: confidentiality, integrity, and availability of information assets, i.e. data or services. This excludes bodily injury and property damage, however, as data is not considered physical property in most cases, includes data loss or data compromise, see [GDV 2017]. First-party coverage typically includes service costs such as forensic investigation costs, costs of crisis management and public relations services, costs of notifying affected individuals (e.g. data breach victims) and reimbursement, as well as BI loss and data restoration loss. Third-party liability coverage covers legal liabilities, e.g. from regulatory defense costs and fines, see [Biener et al. 2015]. While this basic first- and third-party coverage is relevant for all companies seeking stand-alone cyber coverage, additional elements such as coverage of losses from fraudulent use of payment data or damages to Industrial Control Systems might be of interest for some buyers, see [Schonschek 2020].

Many cyber insurance policies seek to transcend mere risk transfer by offering services that help to prevent the manifestation of a cyber incident overall or mitigate the financial consequences of an incident. The incentive for the insurer is to reduce both frequency and severity of losses. Prevention efforts can take the form of software and services included in the policy, e.g. IT security trainings

or special offers on security software. Mitigation efforts include the availability of a 24/7 incident response service, which provides companies with instant advice on how to proceed if they suspect to have fallen victim to a cyber incident. These services are often provided in collaboration with IT security experts. Finally, a concern for many companies is reputation loss as the consequence of a publicly known incident. While the resulting financial loss of a reputation damage is hard or impossible to quantify (and thus to reimburse), cyber insurance coverage can help mitigate the consequences by including legal advice on reporting requirements and crisis communication support, see [Schonschek 2020].

Pricing: (Why) Is my company classified as high risk?
Among the most stated obstacles for purchasing cyber insurance is high premiums, see [Advisen and PartnerRe 2018]. Naturally, companies seek to understand how their risk is classified and priced; while insurers refrain from providing full transparency into their risk assessment and pricing.

In the US, carriers typically assess an applicant's cyber risk through questionnaires, most of which emphasize the amount and type of data handled by the investigated company. Somewhat surprisingly, less attention is put on the technical infrastructure as well as risk management and IT security management. Policies are often priced by multiplying a base premium by variables relating to standard insurance factors (e.g. changes to the limits or deductible and claims history) and industry-related factors, where high hazard weightings are assigned to businesses that collect and store a high volume of particularly sensitive data or operate in industries like healthcare and finance. Finally, premium multipliers are commonly assigned according to the outcome of the IT security questionnaire (e.g. privacy controls, network security controls, and existence of an incident response plan), see [Romanosky et al. 2019]. In contrast to these pricing approaches on the market, contributions from the academic literature on cyber risk modelling seek to study and apply the underlying frequency and severity distributions of cyber losses to risk management and pricing, see e.g. [Eling and Wirfs 2019, Farkas et al. 2020, Zeller and Scherer 2020].

The German Insurance Association's (GDV) sample questionnaire for risk assessment for cyber insurance differentiates between three risk categories primarily according to the annual turnover and secondarily according to certain risky business units (e.g. e-commerce or handling of sensitive data). The number of questions a candidate needs to answer increases with the risk category, starting from 10 general questions regarding policy conditions and 5 questions on specific areas. General questions cover the topics access, protection against malware, patching/updates, and back-up/data storage, whereas additional questions can cover organizational security, separation of networks, and protection of sensitive data, see [GDV 2019] and ▶ Fig. 01. The questionnaire is applicable for small and medium-sized enterprises; for enterprises with annual turnover exceeding 10mn, a more comprehensive risk assessment (e.g. an on-site audit) is advised.

At policy closure, insurers usually require an applicant's IT security standards to meet certain minimum criteria, e.g. usage and regular update of anti-virus software, periodic back-ups with separate storage, patch management, and appropriate access controls such as password protection or two-factor authentication. Duties in case of an incident include immediate notification of the insurer (or its affiliated IT security experts) and support in carrying out the appropriate incident response measures, see [Schonschek 2020].

Investing in cyber security requires strategic decisions and should be done in a comprehensive and effective way. The risk assessment and obligations posed by insurers emphasize that it is not enough to invest in the prevention of cyber incidents. Additionally, the discovery, investigation, and containment of an attack, and the fast recovery of systems to a working state (in accordance with the often-repeated mantra "Assume that you are already compromised") is called for. Studies on the cost of data breaches show that improvements in data governance programs, presence of incident response plans, appointment of a CISO, employee training and awareness programs, and a business continuity management strategy all result in average cost savings in case of a breach, see [Ponemon Institute LLC 2016]. Companies need to accept the reality that cyber risks cannot be fully eliminated and therefore, they must contribute to making the economy more resilient by improving their cyber risk management, including greater investment in security technology and employee awareness.

Accumulation scenarios and the limits of insurability of cyber

In a 2015 survey, many companies stated to not having purchased cyber coverage yet due to unavailability of desired coverage or cover limits, see [Advisen 2015]. This results from a natural conflict of interests: On one hand, buyers of cyber insurance seek coverage for extreme scenarios such as a large public data breach or a long company-wide BI. On the other hand, insurers were approaching this difficult-to-quantify risk by offering (relatively) low cover limits, to be prudent. As the market for dedicated cyber insurance keeps growing and insurers keep expanding their coverage and limits, insurers generally perceive the market as profitable and a strategic growth field.

However, one important concern that leads insurers to tread with caution while expanding their cyber portfolios is accumulation risk. This is caused by interdependence of insureds' systems due to common vulnerabilities (e.g. widespread use of standardized operating systems or the dependence on centralized cloud service providers). Scenarios that are currently regarded as uninsurable due to their underlying accumulation potential and therefore excluded from coverage include the outage of electricity (such as a long-lasting brownout or blackout) and failure within internet infrastructure (e.g. due to a DDoS attack or a telecommunication network outage), see [Munich Re 2020b].

Silent Cyber in Traditional Insurance Lines

Silent cyber describes cover for cyber incidents that may exist in traditional lines, e.g. property/casualty (P&C) policies, even though this was not originally intended by the underwriter. One such scenario is a hacking attack disabling the cooling system of an indus-

Fig. 01: Visualization of the risk assessment questionnaire template by GDV

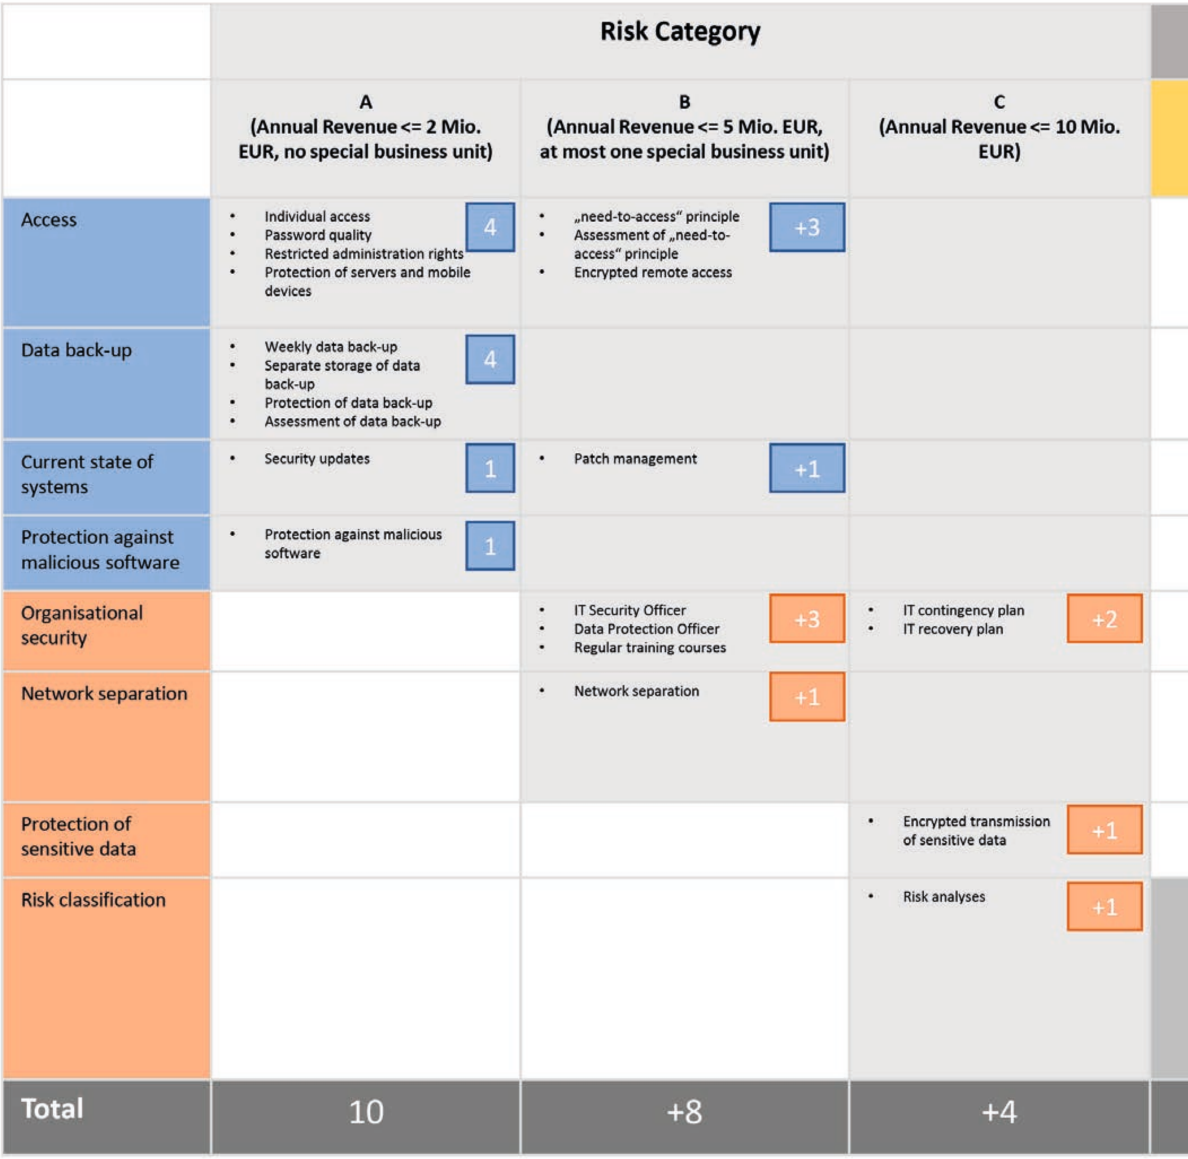

	Risk Category					
	A (Annual Revenue <= 2 Mio. EUR, no special business unit)		B (Annual Revenue <= 5 Mio. EUR, at most one special business unit)		C (Annual Revenue <= 10 Mio. EUR)	
Access	• Individual access • Password quality • Restricted administration rights • Protection of servers and mobile devices	4	• „need-to-access" principle • Assessment of „need-to-access" principle • Encrypted remote access	+3		
Data back-up	• Weekly data back-up • Separate storage of data back-up • Protection of data back-up • Assessment of data back-up	4				
Current state of systems	• Security updates	1	• Patch management	+1		
Protection against malicious software	• Protection against malicious software	1				
Organisational security			• IT Security Officer • Data Protection Officer • Regular training courses	+3	• IT contingency plan • IT recovery plan	+2
Network separation			• Network separation	+1		
Protection of sensitive data					• Encrypted transmission of sensitive data	+1
Risk classification					• Risk analyses	+1
Total	10		+8		+4	

(information from Gesamtverband der Deutschen Versicherungswirtschaft e.V. (2019): Unverbindlicher Muster-Fragebogen zur Risikoerfassung im Rahmen von Cyber-Versicherungen für kleine und mittelständische Unternehmen. (Unverbindliche Bekanntgabe des Gesamtverbandes der Deutschen Versicherungswirtschaft e.V. (GDV) zur fakultativen Verwendung. Abweichende Vereinbarungen sind möglich.)

trial plant, leading to a fire, and resulting in property damage to machines etc. As traditional policies typically describe their coverage along the covered type of loss (e.g. physical damage) instead of its cause, such a scenario would typically be included.

While formerly cyber cover was included as an explicit endorsement to existing policies, as buyers seek dedicated and higher limits and expanded coverage, exclusions in traditional policies become more commonplace and stand-alone cyber products constitute the main source of liability cover. For some coverage elements, such as cyber-related property damage, underwriters now even seem to be divided on whether they should be included in traditional lines or in a stand-alone cyber policy, see [Advisen and PartnerRe 2018]. This inconsistency emphasizes the importance - for insurers and insurance buyers alike – of reviewing existing traditional policies with respect to cyber, to avoid a dangerous perception gap where insureds may suffer from an illusion of protection. For example, the sense of security that any cyber-related BI is covered under a traditional BI policy may be unwarranted if specific exclusions are in place or a physical damage trigger is required.

Summary

As the frequency and severity of cyber incidents is increasing, companies are aware of the imminent need of dealing with cyber as a risk management topic.

While the purchase of external insurance coverage naturally cannot be a replacement for sound internal security measures, novel dedicated cyber insurance products can amend risk assessment and mitigation efforts by providing a multitude of pre- and post-incident services. These services should also be accounted for when the premium of a cyber insurance policy is discussed.

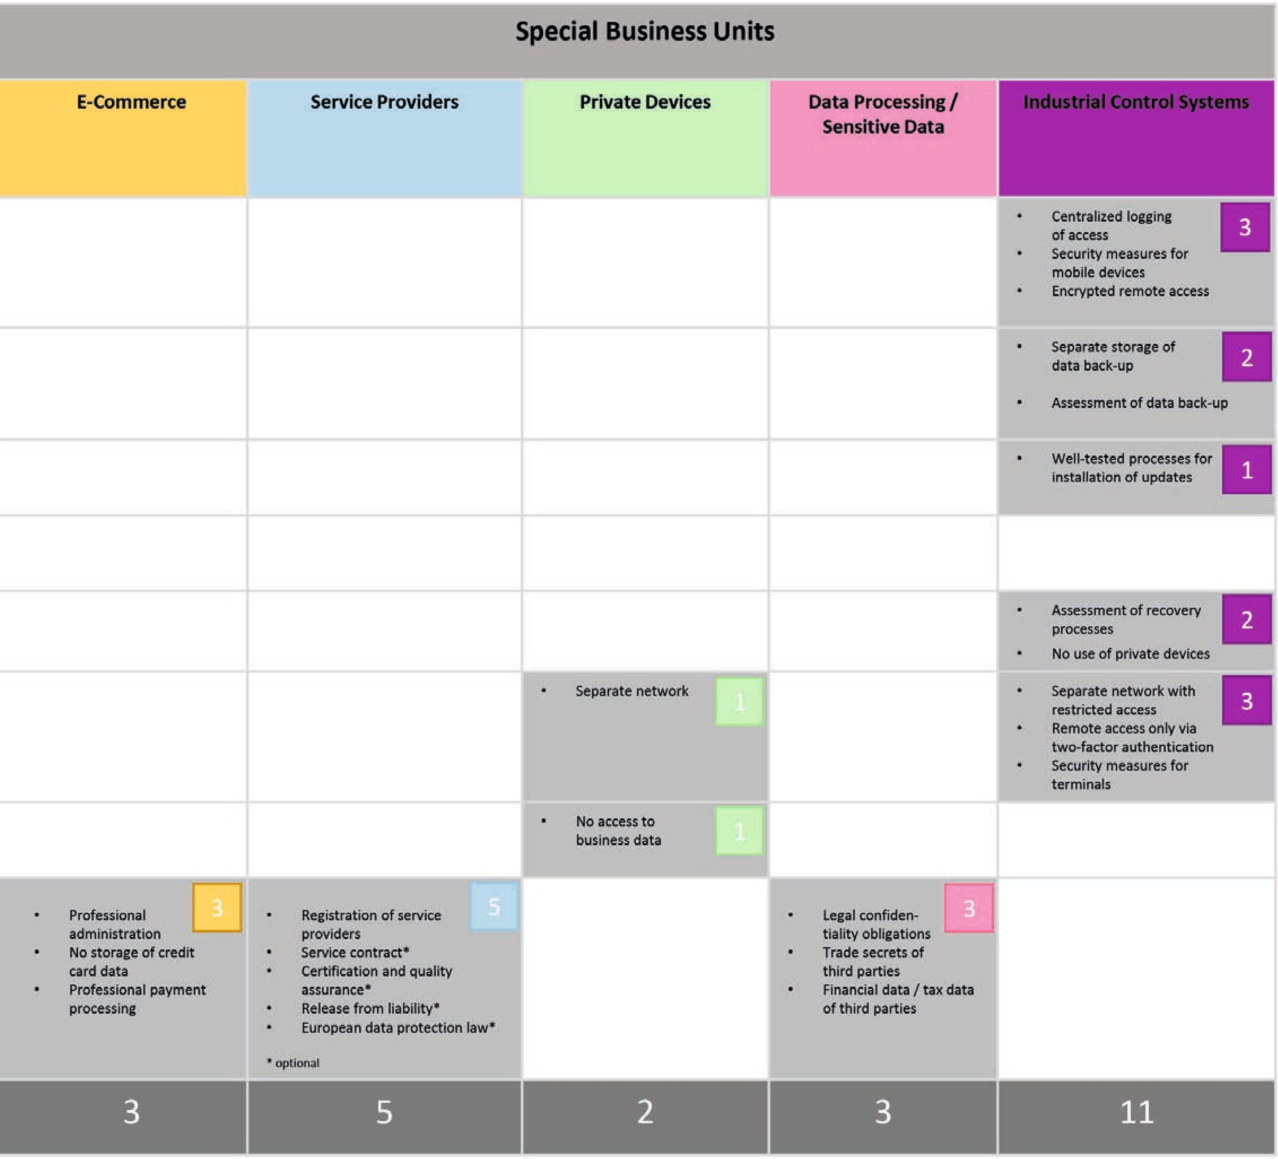

Special Business Units				
E-Commerce	Service Providers	Private Devices	Data Processing / Sensitive Data	Industrial Control Systems
				• Centralized logging of access • Security measures for mobile devices • Encrypted remote access 3
				• Separate storage of data back-up • Assessment of data back-up 2
				• Well-tested processes for installation of updates 1
				• Assessment of recovery processes • No use of private devices 2
		• Separate network 1		• Separate network with restricted access • Remote access only via two-factor authentication • Security measures for terminals 3
		• No access to business data 1		
• Professional administration • No storage of credit card data • Professional payment processing 3	• Registration of service providers • Service contract* • Certification and quality assurance* • Release from liability* • European data protection law* * optional 5		• Legal confidentiality obligations • Trade secrets of third parties • Financial data / tax data of third parties 3	
3	5	2	3	11

Source: own illustration, own translation from German

At first sight, buyers of insurance may find the exact assessment of their coverage needs and their matching with available offers on a rapidly evolving market challenging. The inclusion of cyber-related losses in existing coverage should be critically checked; endorsements to traditional policies may be available, but are predominantly being replaced by stand-alone policies offering higher dedicated coverage and limits as well as invaluable service features.

Companies considering the purchase of dedicated cyber coverage should not let themselves be deterred by the market's lack of transparency, as the decision process itself, including the reflection of one's exposure to cyber threats and the potential (financial) consequences as well as the completion of the insurers' risk assessment process, can certainly serve as the peg on which to hang a potentially overdue conversation about cyber risk management as a central part of ERM – regardless of the final decision to purchase insurance or not. Decision makers need to be aware that in an interconnected world with business models heavily relying on functional internal and external IT infrastructures and digitalization strategies on every board's agenda, complete risk avoidance is not an available strategy for cyber, and the perception that one's own exposure is negligible is almost certainly a fallacy.

Literature

Advisen (2015). 2015 Network Security & Cyber Risk Management: The Fourth Annual Survey of enterprise-wide cyber risk management practices in Europe. (https://www.advisenltd.com/wp-content/uploads/network-security-cyber-riskmanagement-white-paper-2015-02-06.pdf last visited on 19.11.2020)

Advisen and PartnerRe (2017). 2017 Survey of Cyber Insurance Market Trends. (https://partnerre.com/wp-content/uploads/2017/10/PartnerRe-2017-Survey-of-Cyber-Insurance-Market-Trends.pdf last visited on 19.11.2020)

Advisen and PartnerRe (2018). 2018 Survey of Cyber Insurance Market Trends. (https://partnerre.com/wp-content/uploads/2018/10/2018-Survey-of-Cyber-Insurance-Market-Trends.pdf last visited on 19.11.2020)

Advisen (2018). Advisen's 2018 Cyber Guide: The Ultimate Guide to Cyber Service Providers.

Advisen and Zurich (2019). Information Security and Cyber Risk Management. The ninth annual survey on the current state of and trends in information security and cyber risk management.

Allianz Global Corporate & Specialty SE (2020). Allianz Risk Barometer – Identifying the Major Business Risks for 2020.

Allianz Global Corporate & Specialty SE (2019). Allianz Risk Barometer - Top Business Risks for 2019. (https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/Allianz-Risk-Barometer-2019.pdf last visited on 19.11.2020)

Allianz Global Corporate & Specialty SE (2015). A Guide to Cyber Risk: Managing the Impact of Increasing Interconnectivity. (https://www.agcs.allianz.com/content/dam/onemarketing/agcs/agcs/reports/AGCS-Cyberrisk-report.pdf last visited on 19.11.2020)

Biener, C., Eling, M., and Wirfs, J.H. (2015) Insurability of cyber risk: An empirical analysis. The Geneva Papers on Risk and Insurance - Issues and Practice, 40(1), S.131-158.

Eling, M. and Wirfs, J.H. (2019) What are the actual costs of cyber risk events? European Journal of Operational Research, 272(3):1109-1119.

Farkas, S., Lopez, O., and Thomas, M. (2020) Cyber claim analysis through Generalized Pareto Regression Trees with applications to insurance pricing and reserving. https://hal.archives-ouvertes.fr/hal-02118080 .

Gesamtverband der Deutschen Versicherungswirtschaft e.V. (2019). Unverbindlicher Muster-Fragebogen zur Risikoerfassung im Rahmen von Cyber-Versicherungen für kleine und mittelständische Unternehmen. (Unverbindliche Bekanntgabe des Gesamtverbandes der Deutschen Versicherungswirtschaft e.V. (GDV) zur fakultativen Verwendung. Abweichende Vereinbarungen sind möglich.) (Free download available: https://www.gdv.de/resource/blob/6102/aa4b9afe6fa3e23c51c598bd23194ba1/02-risikofragebogen-cyber-data.pdf last visited on 19.11.2020)

Gesamtverband der Deutschen Versicherungswirtschaft e.V. (2017). Allgemeine Versicherungsbedingungen für die Cyberrisiko-Versicherung (AVB Cyber). Musterbedingungen des GDV. (Unverbindliche Bekanntgabe des Gesamtverbandes der Deutschen Versicherungswirtschaft e.V. (GDV) zur fakultativen Verwendung. Abweichende Vereinbarungen sind möglich.) (Free download available: https://www.gdv.de/resource/blob/6100/d4c013232e8b0a5722b7655b8c0cc207/01-allgemeine-versicherungsbedingungen-fuer-die-cyberrisiko-versicherung--avb-cyber--data.pdf last visited on 19.11.2020)

Munich Re (2020a). Cyber insurance: Risks and trends 2020. (https://www.munichre.com/topics-online/en/digitalisation/cyber/cyber-insurance-risks-and-trends-2020.html last visited on 19.11.2020)

Munich Re (2020b). What if a major cyber attack strikes critical infrastructure? (https://www.munichre.com/topics-online/en/digitalisation/cyber/silent-cyber.html last visited on 19.11.2020)

Ponemon Institute LLC (2016). 2016 Cost of Data Breach Study: Global Analysis. (https://www.academia.edu/35179110/2016_Cost_of_Data_Breach_Study_Global_Analysis last visited on 19.11.2020)

Romanosky, S., Ablon, L., Kuehn, A., and Jones, T. (2019). Content analysis of cyber insurance policies: how do carriers price cyber risk? Journal of Cybersecurity, 5(1), S.1-19.

Schonschek, O. (2020). Letzte Rettung. Marktübersicht: Cyberrisiken versichern lassen. iX 4/2020, S.48-62.

Zeller, G. and Scherer, M. (2020). A comprehensive model for cyber risk based on marked point processes and its application to insurance. Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3668228

Authors

Gabriela Zeller

Doctoral Candidate,
Chair of Mathematical Finance,
Technical University of Munich

Prof. Dr. Matthias Scherer

Professor for Risk and Insurance,
Chair of Mathematical Finance,
Technical University of Munich

[Image source: Astrid Eckert]