CAS RESEARCH PAPER

CYBER RISK: QUANTIFICATION, STRESS SCENARIOS, MITIGATION, AND INSURANCE

Olivier Lopez, Michel Denuit, Mario Ghossoub, Julien Trufin, Justin Kher, Arthur Maillart, Elisabeth Raes, Hugo Rapior, Mohammed-Amine Skoubani, Brieuc Spoorenberg

CASUALTY ACTUARIAL SOCIETY

The Casualty Actuarial Society (CAS) is a leading international organization for credentialing, professional education and research. Founded in 1914, the CAS is the world's only actuarial organization focused exclusively on property-casualty risks and serves over 10,000 members worldwide. CAS members are sought after globally for their insights and ability to apply analytics to solve insurance and risk management problems. As the world's premier P&C actuarial research organization, the CAS reaches practicing actuaries across the globe with thought-leading concepts and solutions. The CAS has been conducting research since its inception. Today, the CAS provides thousands of open-source research papers, including its prestigious publication, Variance — all of which advance actuarial science and enhance the P&C insurance industry. Learn more at casact.org.

Caveat and Disclaimer

This research paper is published by the Casualty Actuarial Society (CAS) and contains information from various sources. The study is for informational purposes only and should not be construed as professional or financial advice. The CAS does not recommend or endorse any particular use of the information provided in this study. The CAS makes no warranty, express or implied, or representation whatsoever and assumes no liability in connection with the use or misuse of this study. The views expressed here are the views of the authors and not necessarily the views of their current or former employers.

The views reflected in this CAS Research Paper are the views of the authors and do not necessarily reflect the views of the global PwC organization or its member firms. The information contained in this publication is of a general nature only. It is not meant to be comprehensive and does not constitute the rendering of legal, tax or other professional advice or service by PwC China or PwC Hong Kong. PwC has no obligation to update the information as law and practices change. The application and impact of laws can vary widely based on the specific facts involved. Before taking any action, please ensure that you obtain advice specific to your circumstances from your usual PwC client service team or your other advisers. The materials contained in this publication were assembled on August 18, 2024, and were based on the law enforceable and information available at that time.

Contents

Cyber Risk: Quantification, Stress Scenarios, Mitigation, and Insurance

Olivier Lopez,[1,2] Michel Denuit,[3,4] Mario Ghossoub,[5]
Julien Trufin,[3,6] Justin Kher,[1] Arthur Maillart,[1] Elisabeth Raes,[3]
Hugo Rapior,[1] Mohammed-Amine Skoubani,[1] Brieuc Spoorenberg[3]

The paper discusses challenges that arise in the quantification of cyber risk in the context of cyber insurance. We propose a methodology for developing a synthetic database of cyber events that may serve as a benchmark for risk analysis. The database is designed from the perspective of anticipating and optimizing risk transfer and risk management procedures. It can also be used to test the impact of different prevention schemes. Last, we discuss the question of stress-testing insurance portfolios by analyzing how to calibrate different stress scenarios that may lead to a failure of mutualization. All the methodologies developed herein are presented in such a way that one can quickly adapt and update them in light of new information and expertise so as to take into account the evolution of risk.

Keywords: Cyber risk, Cyber insurance, Generalized linear mixed models, Mixed Poisson models, Cyber risk stress scenarios

1. Introduction

An earlier CAS research paper on cyber risk insurance (Bean 2020) provided an overview of the different ways by which one can measure exposure to cyber risk. It highlighted the variety of cyber risk events and attacks, as well as some of the risk factors that tend to increase or reduce exposure and that can be used for pricing or prevention. The present proposal aims to supplement that first analysis with technical methods the actuary can use to go beyond the identification of risk factors and quantitatively evaluate their economic impact on insurance contracts. Clearly, measuring the impact on policyholders of risk factors or characteristics is part of standard actuarial methodologies. Nevertheless, the context of cyber risk is quite specific, given its novelty, its fast evolution, and the significant amount of the risk exposures typically at stake.

[1] Detralytics, Paris, France.
[2] CREST Laboratory, CNRS, Groupe des Écoles Nationales d'Économie et Statistique, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France.
[3] Detralytics, Brussels, Belgium.
[4] Institute of Statistics, Biostatistics and Actuarial Science - ISBA, Louvain Institute of Data Analysis and Modeling - LIDAM, UCLouvain, Louvain-la-Neuve, Belgium.
[5] Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, Ontario, Canada.
[6] Department of Mathematics, Université Libre de Bruxelles (ULB), Brussels, Belgium.

In this paper, we investigate how to adapt the actuarial toolbox to these specificities, with a particular focus on the combination of historical data and expert judgment, as well as on how to use these sources of information in order to build risk transfer solutions. To that end, we explore a few main directions in the study of cyber risk and cyber risk insurance, namely:

1. *How to couple classical Bayesian analysis with relevant expert judgment in cyber risk modeling*. Given the lack of reliable cyber risk data, expert judgment often becomes a crucial part of the risk quantification and measurement process; however, expert judgment may be hard to combine with established actuarial methodologies. For instance, there might be a discrepancy between expert judgment and the scientific requirements of statistical models, and/or the parameters of a specific cyber insurance contract. This is especially true for official reports from security administrations that provide a well-documented analysis of the threat but are not directly suited for insurance applications. This paper suggests potential solutions for reconciling these two levels of information of different types into a more accurate evaluation of the risk.
2. *How to account for various forms of heterogeneity in claims, data, impact, and so forth in the evaluation of cyber risk, as well as in the pricing and/or design of risk transfer mechanisms*. Cyber claims can be of various types, and their consequences may be quite different from one type of policyholder to another. Additionally, the parameters of cyber insurance contracts may vary (e.g., Romanosky et al. 2019). It is therefore crucial to account for this and other forms of heterogeneity in the evaluation of cyber risks, as well as in the pricing and/or design of risk transfer mechanisms. Traditional methods focus on the impact of risk factors on a given "central" scenario. For instance, generalized linear models mainly focus on the mean. Here, we leverage recent techniques (e.g., Farkas et al. 2021) that can help the actuary identify and classify policyholders and/or cyber events with respect to the extreme (catastrophic) scenarios that could arise.
3. *How to use risk transfer solutions in the management of cyber risk*. Risk transfer solutions are crucial in cyber risk management because of the significant potential claim severity. The viability of such solutions, which will be addressed in this paper, relies essentially on anticipating extreme scenarios (that is, analyzing the tail of the distribution of the losses), as well as achieving a sufficient size for the exposed portfolio. The methods mentioned in the previous point to account for heterogeneity will help explain how to design risk transfer strategies that adapt to the particular shape of the portfolio: categories of policyholders/contracts with lesser catastrophic risk could generate more cost-efficient risk transfer strategies that may contribute to attracting new policyholders and improving the effects of pooling. However, quantitative methods may help to determine which type of claims cannot be covered and to establish a frontier between insurability and non-insurability.
4. *How to effectively develop stress-testing scenarios*. Apart from the classical frequency-severity approach, a major concern regarding cyber risk is the fear of a "cyber hurricane" that would generate a failure of a large part of a given portfolio. We propose tools for building stochastic scenarios that can be used to evaluate the robustness of risk transfer strategies. The description of these tools is supplemented

by an analysis of the main deterministic catastrophic scenarios the sector might face (e.g., a cyberattack on a cloud provider).

5. *How prevention can be used as a risk mitigation or risk management strategy*. Prevention can be seen as a particular risk transfer strategy that presents the advantage of promoting the role of insurers as major actors in the cybersecurity community. We provide and illustrate methods of quantifying the impact of prevention and its associated financial cost, thereby yielding an objective view of how prevention may be profitable to the sustainability of a given portfolio.

Additionally, and as mentioned above, cyber risk analysis suffers from a lack of public data that would otherwise allow for a proper risk evaluation and calibration of actuarial models. A key contribution of this paper is the construction of a dataset reflecting the main figures in the market, based on our methodology that effectively integrates expert judgment. This artificial dataset is made of simulated claims enabling the illustration of the methodologies we use and the benchmarking of future actuarial models. The parameters used to generate such a dataset are tractable – a user may modify them to adapt to changes in the environment. This question of adaptation to an evolving environment is crucial, precisely because of the quickly changing landscape of the cyber risk environment. Although any dataset would soon be obsolete due to these changes, the procedure that we propose presents the advantage of providing reasonable and publicly available datasets for the community that can be easily updated.

The rest of this paper is organized as follows. Section 2 provides a broad overview of the main complications one encounters when defining cyber risk, classifying its different types, modeling it, and quantifying it. Section 3 presents the main mathematical framework used in modeling the frequency of cyber risk events. Section 4 describes the methodology used to generate the database of cyber claims, introduces a taxonomy of cyber events, discusses the procedure used to model the severity of events, and illustrates how the proposed methodology can be used to generate portfolios with different structures. In Section 5, we discuss potential risk transfer solutions and parametric cyber risk insurance products, as well as the role that prevention can play in cyber risk management. Section 5 also discusses how our proposed framework can be used to generate stress scenarios for cyber risk. Background material can be found in the appendix.

2. Difficulties in Quantifying Cyber Risk

2.1. The Perimeter of Cyber Risk

The concept of cyber risk covers an important variety of situations. In full generality, *cyber risk* refers to the risk of a failure in information systems that leads to damages. On the other hand, cyber insurance products are mostly focused on "malicious cyber," that is, cyber triggered by attacks from criminal groups. Romanosky et al.'s (2019) study on the content of various policies on the US market shows that the main concern of this field of insurance is to respond to malicious acts targeted against a company. The non-malicious part of cyber is typically linked to traditional industrial liability insurance and is therefore beyond the scope of the present paper. However, it is important to keep this component in mind

as the consequences of some purely accidental incidents can give us clues about the consequences of potential cyberattacks (an example is a cloud outage, such as that experienced by Amazon – see Hagen et al. [2012]; see also Li et al. [2013]).

The question of the definition of cyber is also key in the context of some catastrophic events that cannot be absorbed by the insurance market. The case of *cyber war* is an obvious example, but the definition of cyber war itself is not clear. During the NotPetya incident, which was explicitly attributed to a foreign power by several intelligence services, some cyber insurers attempted not to cover the damages because of war exclusion (Wan 2020), without success. Recently, Lloyd's of London asked their syndicates to include a clause of exclusion of "state-sponsored attacks," but the adaptation of such clauses to other legislation may be considered with caution.

Although, according to Romanosky et al. (2019), the content of cyber policies tends to be relatively homogeneous, the options offered to customers can vary, making it difficult to compare the severity of a given cyber event from the sole perspective of the amount of insurance compensation received by the victim.

The particular case of ransom payment is a good example: on this sensitive point, authorities strongly recommend not including this option systematically in the policy or including terms that avoid, as often as possible, triggering this part of the guarantee. The consequence of systematizing the ransom payment through insurance contracts would be to put a target on the policyholder.

2.2. Some Typical Cyberattacks

The following is a nonexhaustive list of classic types of cyberattacks and of ways to introduce a virus into an information system. One can begin to see the heterogeneity of situations behind cyber, in terms of patterns of attack, timeline, and consequences.

- Distributed denial of service (DDoS): A huge number of requests is sent to a server, which becomes paralyzed. The requests are usually sent by botnets, that is, viruses inside infected computers or connected objects around the world that are activated at the proper moment to trigger the DDoS attack. The aim of this type of attack is to freeze some target during a (usually) short but strategic period of time.
- Credential theft: A hacker enters the system using a user's password. Hackers use various techniques to steal passwords, from stealing a piece of paper on which a password was imprudently written to "brute force attacks," where a program repeatedly guesses passwords until it finds one (usually because that password was too simple).
- Phishing: The hacker uses a fraudulent email to try to convince a worker to reveal their password (in which case, phishing is a gateway to credential theft) or to download malware that infects the system.
- Insider: Someone inside the company voluntarily introduces the virus (e.g., using an infected USB key) directly into the targeted computers.

- Ransomware: The virus blocks access to the information system unless a ransom is paid in a given amount of time (note that payment does not always allow complete retrieval of the affected data). A "wiper" can destroy the data following the expiry of the deadline, and/or the hackers may blackmail the victim, threatening to disclose some sensitive information they have accessed (this is called "double extortion"). Business interruption is also an important consequence of such an attack.
- Data breach: Data are exposed and/or stolen. Disclosure is not automatic. Regulations pertaining to data privacy can lead to financial penalties. Reputation can also be damaged.
- CEO fraud: Social engineering methods are used to fool unwitting collaborators by impersonating a person with high authority – to obtain, for example, an illegal transfer of funds. This technique usually requires access to confidential information (e.g., via hacking) in order to make the deception more plausible.
- SQL injection attack: The hacker intercepts an SQL request – for example, to introduce false information into databases or to mislead the victim.
- Cryptojacking: A virus diverts computing capacity of the victim to mine cryptocurrencies. The victim may be unaware of the infection but may suffer excessive energy costs and loss of efficiency of internal processes.

Again, the list is not exhaustive, but it demonstrates that cyber may take many forms and can lead to various consequences that may not always be immaterial, such as in the case of business interruption.

2.3. Lack of Data

In his CAS research report, Bean (2020) discusses several potentially valuable sources of data for use in analyzing cyber risk. However, by design, those sources of information are not perfectly fitted to insurance pricing or reserving, the main reason being that, historically, the question of cyber is of a technological nature, and therefore most of the databases collecting cyber incidents are designed by computer scientists for the sole purpose of contributing to improving the safety of information systems. For example, the National Institute of Standards and Technology's National Vulnerability Database (see https://nvd.nist.gov/) lists vulnerabilities detected by the cybersecurity community, but the consequences of an exploit based on a given vulnerability are hard to determine. In particular, the economic consequences of an attack are hard to match up with a given vulnerability. The case of the Log4shell vulnerability (Everson et al. 2022) is an example of a widely shared unanticipated weakness, with no precise clue about the consequences.

Incidences of data breaches, compared with other kinds of cyber risk, are relatively well documented. The initial importance accorded to data breaches relates to the attention paid to the issue of respect of privacy: one of the most-used databases in the field of research on cyber insurance is maintained by Privacy Rights Clearinghouse (PRC; see https://privacyrights.org/data-breaches), but the initial purpose of the clearinghouse and its database, to supply information about protecting US citizens' data privacy rights,

is a far cry from addressing the question of designing insurance contracts. For risk evaluation, the case of data breaches is interesting, because it provides an element of severity, namely, the volume of exposed data. But the link between that volume and the corresponding economic loss is hard to properly quantify. The scientific literature often relies on an old formula introduced by Jacobs (2014) based on data from the Ponemon Institute. The reliability of such a formula is questionable: Farkas et al. (2021) showed that it was not suited to more recent "mega-breaches" (such as the Yahoo data breach in 2016 – see Thielman [2016]), and proposed a modification, which seems better adapted but lacks guarantees regarding the certainty of such a formula. On the other hand, some studies are specifically dedicated to the question of cyber insurance and may help us begin to picture the risk. Here we focus on two of those that seem to us the most reliable (although none is exempt from weaknesses).

The LUCY (*Light Upon CYber insurance*) study (see AMRAE 2023), conducted by the French association AMRAE since 2021 (with data since 2019), provides an exhaustive view of losses on the French cyber insurance brokers' market. The available information is relatively scarce, however, as the study gives only a distribution of the losses on a relatively small number of segments. Another important drawback is the study's bias. By focusing on a single market (France) and on policyholders relying on brokers, LUCY misses some important information. Although the study may provide some interesting elements about the magnitude of claims striking large insured companies, it fails to capture most of the information regarding smaller actors, such as small- to medium-sized enterprises, which rarely rely on brokers. But the main benefit of the approach is to clearly define its perimeter: we know precisely how the sample has been constituted, and hence we know the exposure. That advantage is missing from and constitutes the major drawback of the next study we focus on.

The annual Hiscox Cyber Readiness Report (see Hiscox 2022) is a valuable source of information regarding the evolution of cyber risk. We will consider it in this paper as a benchmark for calibrating the reference models that we elaborate. The report is based on an inquiry sent to various companies operating in different sectors in eight countries. It collects information about the nature of those companies' cyber activities and events and their consequences. Even if the number of countries is relatively limited, and even if some of the responses are aggregated to respect privacy, the study has the significant advantage of providing information about most of the questions that a cyber insurer would like to ask to evaluate risk. However, it is important to acknowledge the report's limits, with the main one being a lack of transparency about the way the sample has been determined. Its representativeness is therefore unclear.

Therefore, although it constitutes a plausible reference that is probably, for now, one of the most reliable sources of information, the Hiscox report can be considered as biased. Hence, when dealing with an insurance portfolio, it is crucial to progressively eliminate that bias by gathering more experience. We describe a possible way to proceed in Section 3.

2.4. Heavy-Tailed Losses

One difficulty in risk analysis in the context of cyber is the extreme volatility of the loss. We can observe that extreme volatility, for example, in the LUCY study. According to AMRAE's survey, the loss ratio in the French market (restricted to contracts sold by brokers)

is evaluated at 167%. But if we remove just 4 of the 182 claims the study considers, the ratio falls to 85%.

From a statistical point of view, stylized facts corroborate this important disparity between the potential losses if we consider the well-documented example of data breaches, keeping in mind that data breaches make up only a small part of cyber risk. Indeed, Maillart and Sornette (2010) documented, from an analysis of the PRC database, that a heavy-tailed distribution (a "power law distribution") described the volume of leaked data. Their finding was explored further by Edwards et al. (2016) and Farkas et al. (2021), who, with more recent data, arrived at an even more pessimistic analysis regarding the heaviness of the tail.

Technically speaking, a Pareto tail means that if we consider the random loss Y of a cyber event (for a single policyholder), its distribution is of the type

$$S_Y(y) = \mathbb{P}(Y \geq y) = \frac{l(y)}{y^{1/\gamma}},$$

where $l(y)$ is a slowly varying function[1] and $\gamma > 0$ is the tail index. A rigorous definition of this class of distributions and their importance in extreme value theory can be found, for example, in Beirlant et al. (2004). The key idea behind this mathematical definition is that the decrease of S_Y is slow compared with more convenient variables, such as with Gaussian, gamma, or even lognormal distributions. We can draw a couple of worrisome conclusions from this:

- It is not rare to experience scenarios that are far from the "average," that is, from the expectation of the variable (used in the computation of the pure premium in insurance), since the probability of having Y larger than some large amount is not so small.
- At a portfolio level, if we consider the total loss of the insurer, the Gaussian approximation of this loss may not be appropriate, especially when one focuses on high quantiles – see, for example, Mikosch and Nagaev (1998). This could be concerning, because such an approximation is regularly used in computing reserves; for example, it is behind the standard formula of the European Solvency II directive.

The problem can be even worse if γ is too large, as in the case of particularly volatile losses. If $\gamma > 1$, the expectation itself is infinite, making it impossible to define a pure premium. This may seem unrealistic as the loss of a cyber incident may be huge, but it is not infinite. It simply means that, at some point, the risk may be so volatile that its mathematical modeling reaches its limits. In this case, securing the tail of the distribution by gathering a sufficiently large reserve is hopeless. Even risk transfer to reinsurance will not enable the insurer to cover the risk entirely because the reinsurance premium cannot even be computed. If $\gamma > 0.5$ the risk is insurable in the mathematical sense of the term, but the variance of the loss becomes infinite, which signifies that the traditional risk mitigation strategy based on return/variance compromise is no longer possible.

[1] Essentially this means that $l(y)$ can be considered as constant for y large (or with a logarithm type increase that is negligible compared to the denominator $y^{1/\gamma}$).

In any case, the power law decay is a warning sign, even for small values of γ, indicating that managing the risk management will necessitate particular caution.

This discussion may seem, at first glance, essentially theoretical. A simple solution exists, via the introduction of a limit in the compensation. If the contract says that the insurance company does not pay more than a given limit M, then, from the insurer's point of view, the paid loss (say L) is bounded. In such a case, there would apparently be no need to concern ourselves with the tail of the distribution and with extreme behavior that may go beyond the scope of compensation. However, this ignores at least two aspects of the problem:

- First, to fix the boundary of compensation M, one must understand the true loss Y experienced by the policyholder. From the knowledge on the tail, we will be able to fix the appropriate threshold, and we will be able to quantify how frequently this limit of compensation is reached.
- Second, the insurance company needs to cover a significant part of the risk. If not, the policyholder may be reluctant to subscribe, hence slowing the growth of the portfolio (and therefore making mutualization more difficult to achieve). This potential unappealing aspect of cyber insurance contracts is difficult to properly quantify. Nevertheless, AMRAE's LUCY report on the French market in 2022 (AMRAE 2023) measured such shrinkage in some segments of the market, mentioning the strategy of large groups preferring to spend their budget directly on cybersecurity due to the lack of proper coverage. Although the current growth of the cyber insurance market (according to Fitch, a 50% increase in written premiums in 2022, following 73% in the previous year) may seem reassuring regarding this issue, the relative youth of the market invites caution; reversals linked to the evolution of the threat may disrupt this equilibrium.

For these reasons, in the following we choose to focus on modeling the effective total loss experienced by the policyholder rather than on the insured loss itself. The loss may indeed vary depending on the conditions of the policy. Moreover, focusing on the true loss allows us to model more freely the rules of compensation and the design of the contract itself.

2.5. Accumulation Risk

Let us link the previous discussion with a traditional frequency-severity approach. That traditional framework assumes that there is, more or less, independence among policyholders – that is, two policyholders will rarely be struck at the same time by a cyberattack. The equilibrium of the portfolio benefits from the fact that the important losses suffered by a small number of policyholders are absorbed by the premiums paid by the vast majority of policyholders who experience no claim.

On the other hand, for insurers of cyber risk the possibility of an accumulation event – that is, a catastrophe that strikes a large number of policyholders in a short amount of time – is a major concern. This concern stems from the viral nature of some cyberattacks. The virus may spread within a company's information system, but it can also spread via internet connections. Numerous examples of such episodes can be found in the history

of cybersecurity, up to now with relatively moderate impacts on the insurance sector. Here are a few examples:

1. In May 2000, during the internet's early days, the "Love Bug" generated an estimated US$5.5–8.7 billion in damages worldwide. The worm scanned the address books of victims, sending them infected emails.
2. In one week in May 2017, the WannaCry ransomware attack struck around 300,000 computers in 150 countries, causing an estimated loss of US$2 billion. The attack exploited a Microsoft operating system vulnerability that was well known but had not been corrected for many users.
3. The NotPetya attack was similar to the WannaCry attack in terms of point of entry (same vulnerability as WannaCry), but the pattern was different, as some parties called it an act of "cyber war" (Woods and Weinkle 2020; Wolff 2021). The presence of Russian language in the code and the fact that the attack initially targeted Ukraine may corroborate this diagnosis. However, Zurich Re's attempt to refuse compensation to Mondelez, claiming that war (hence cyber war) was excluded from the policy led to a legal dispute ending with the victory of the policyholder. This example shows the difficulty insurers face in properly qualifying, and potentially excluding, state-sponsored attacks from coverage.
4. Other cases of so-called "cyber war" have been documented. One may list the attack against Australia during the COVID-19 crisis, considered as a retaliation against political positions of the country (Lallie et al. 2021). In such cases, the impact on the insurance sector (and its role, with the potential exclusion of such episodes) is unclear.

Hitherto, the impact of such episodes on insurance has been relatively moderate. However, one can argue that the magnitude of these crises is low compared to what one could expect from a huge cyber catastrophe. Recent modeling, for example, of the dynamics of the WannaCry episode (Hillairet and Lopez 2021) showed that the contagiousness of the attack was low. If you were to compare it to a biological virus, its rate of replication, which drives the size of the epidemic, would be only a few digits above the minimum value required for the crisis to spread. That is why the European Insurance Occupational Pensions Authority recently emphasized the need for calibrating stress scenarios that could quantify "cyber underwriting risk" – that is, the risk of massive failure of a portfolio related to a global cyber event.

A massive failure could be produced by a means other than a purely contagious process, however. The cloud outage scenario is one of the more concerning ones (Reece et al. [2023] – see also the report of the European Insurance and Occupational Pensions Authority [EIOPA 2023] and the US Treasury Department's report on the specific case of the financial sector at https://home.treasury.gov/system/files/136/Treasury-Cloud-Report.pdf). An attack on or spreading through a cloud provider could, because of the concentration of the sector into a small number of actors, lead to huge damages. From a modeling perspective, the consequences of such a scenario would be easier to identify than those of the previously mentioned contagion events as long as the insurance company kept track of the cloud solutions used by its different customers.

However, whereas cloud provider failures are clearly identified threats, silent threats are also a matter of concern. Indeed, policyholders share many digital solutions, apart from cloud solutions. The Log4shell vulnerability (Everson et al. 2022) is a recent example of such a threat – an unexpected vulnerability that affects a large number of users (the exact number not being known) and that suddenly appears with no clear anticipation of its consequences.

In Section 5.4 we discuss how to approach the problem of evaluating the impact of an accumulation scenario such as those just mentioned (contagious, cloud outage, or failure of a shared digital solution).

3. A Generalized Linear Mixed Model Framework

3.1. Mixed Models

Datasets exhibiting a hierarchical or nested structure or including longitudinal or spatial elements often arise in insurance studies. This generally results in correlation among the responses within the same group. Random effects offer a convenient way to model such a grouping structure. Here we briefly introduce the generalized mixed model, or GMM, approach to regression analysis. In this framework, available features (called fixed effects) are combined into a score and random effects are added on the score scale. With generalized linear mixed models, or GLMMs, the score is a linear combination of the features and random effects, and the expected response is mapped on the score scale by a monotonic link function. Predictive distributions, that is, a conditional distribution of the response given past experience, are particularly attractive to reevaluate future premiums based on claims observed previously, bridging mixed models to actuarial credibility theory. We refer the interested reader to Ohlsson and Johansson (2010) for more about GLMMs.

Mixed models retain most of the structure of standard regression models except that the score now involves a random part. The features $x_{i,j}$ recorded in the data basis for policyholder i are referred to as fixed effects, supplemented with random effects generating the correlation structure. With GLMMs, for instance, the score η_i is equal to $\sum_j \beta_j x_{i,j} + \Delta_i$, where the linear combination $\sum_j \beta_j x_{i,j}$ corresponds to the GLM score and Δ_i represents the random effect. Contrary to the fixed effects, which are known to the analyst, random effects are random variables accounting for the missing information, i.e., for the heterogeneity that is present in the portfolio beyond that captured by the available features $x_{i,j}$. Random effects are generally structured in a hierarchical or other meaningful way.

In the statistical literature, GLMMs correspond to a wide class of model, where the expected response μ_i for policyholder i is mapped to the score scale by some link function g, that is, $g(\mu_i) = \eta_i$, where

$$\eta_i = \beta_0 + \sum_j \beta_j x_{i,j} + \sum_k B_k z_{i,k},$$

so that $\Delta_i = \sum_k B_k z_{i,k}$. Here, the vector $(B_1, B_2, \ldots)$ of random effects is multivariate normal with zero mean and a structured variance-covariance matrix to be estimated from data. This means

that the regression coefficients corresponding to the features $z_{i,k}$ are no longer constant but become random.

The likelihood can be obtained by integrating out the random effect. Because of the presence of the integrals, maximum likelihood estimation requires numerical procedures. The maximization of the likelihood can be performed, for instance, in one of the two following ways:

- By approximating the likelihood using Laplace's method. This leads to a quasi-likelihood augmented with a penalty term. This first approach is referred to as the penalized quasi-likelihood, for which a version of the Fisher scoring algorithm is available.
- By approximating the integral using numerical integration techniques. Adaptive Gauss-Hermite quadrature formulas can be used to evaluate the integral involved in the likelihood (replacing it with a weighted sum).

As an alternative, generalized estimating equations (GEEs) account for the dependence present in the data in a relatively easy way. GEEs provide a practical method with reasonable statistical efficiency to analyze correlated data.

3.2. A Mixed Poisson Model for Claim Frequency

Consider an insurance portfolio comprising n policies. Let N_i be the number of claims reported by policyholder $i \in \{1, 2, \ldots, n\}$ during the observation period. At the beginning of the coverage period, the actuary has at their disposal some information about each policyholder summarized into a set of features $x_{i,j}$ gathered in the vector $\boldsymbol{x}_i$. The a priori information $\boldsymbol{x}_i$ is recorded in the data basis under consideration. Resorting to regression (or supervised learning) machinery recalled earlier, this information is integrated into the prediction of the annual expected number of claims, or claim frequency. Henceforth, d_i refers to some meaningful exposure to risk for policyholder i, measured by number of employees or number of PCs, for instance.

As explained earlier, a random effect Δ_i is included in the analysis. The numbers of claims N_i are then assumed to be independent given Δ_i. The latent, hidden random effects Δ_i characterize the correlation structure of the claim counts N_i. These random effects can be structured in a number of meaningful ways. For instance, Δ_i can be decomposed into the sum of a shared country effect $\Delta_{c,i}$ common to all policies written in the same country c, and a shared sector effect $\Delta_{s,i}$ common to all firms active in the same economic sector s, as in the next sections. The model is based on the following assumptions. First, given $\Delta_i = \delta$, the random variables N_i are independent and conform to the Poisson distribution with mean

$$\lambda_i \exp\left(\delta\right) = \exp\left(\ln d_i + \eta_i + \delta\right).$$

Adding $\ln d_i$ to the score (i.e., treating this quantity as an offset) means that the insurer's price list is expressed per unit of exposure and varies according to traditional risk features included in the vector $\boldsymbol{x}_i$. Second, the random effects Δ_i obey the multivariate normal distribution with zero mean and variance-covariance matrix $\sum_{\Delta}$.

When the canonical log link function is used in the Poisson regression model, as assumed here, this amounts to using a multivariate Poisson-lognormal model for claim counts. Contrary to what is generally assumed in the actuarial literature, where the random effects Δ_i are supposed to be such that $E[\exp(\Delta_i)] = 1$, the statistical literature devoted to mixed models assumes that the random effects Δ_i are centered. We then have $E[\exp(\Delta_i)] = \exp(\mathrm{Var}[\Delta_i]/2)$ according to the formula giving the mathematical expectation for the lognormal distribution. Hence,

$$E[N_i] = E\left[E[N_i|\Delta_i]\right] = \lambda_i E\left[\exp(\Delta_i)\right] = \lambda_i \exp\left(\mathrm{Var}[\Delta_i]/2\right).$$

The nature of the dependence induced by the mixed models is appropriate for cyber risks. Contrary to common shock models where entire blocks of policies are hit by the same event (as in natural catastrophes), cyberattacks here rather tend to increase the claim frequencies for policyholders sharing the same random effects (without systematically resulting in additional claims).

Mixed models used in statistics are intimately connected to actuarial credibility models for experience rating. Experience rating consists of integrating past claims experience into next year's premium. To link the mixed-modeling approach to credibility, assume that panel data are available. This means that policyholder i is followed over time, say for T_i periods. Compared with classical actuarial studies that deal with annual periods, shorter time periods, such as a quarter or a month, may be appropriate for a rapidly evolving risk like cyber that needs close monitoring and may exhibit some seasonality. Let $N_{i,t}$ be the number of claims reported by policyholder i, $i = 1, 2, \ldots, n$, during period t, $t = 1, 2, \ldots, T_i$. The random effect Δ_i can be decomposed into shared random effects supplemented with an individual one, say E_i. The distribution of E_i can be revised in a Bayesian way based on past claim experience $N_{i,1}, N_{i,2}, \ldots, N_{i,T_i}$ to obtain the predictive distribution of N_{i,T_i+1} given $N_{i,1}, N_{i,2}, \ldots, N_{i,T_i}$. Linear credibility formulas can also be used to that end.

3.3. Bayesian Framework and Expert Judgment

The Bayesian approach is designed to deal with situations where the amount of data is insufficient to rely on a purely frequentist approach. This general framework takes into account an expert judgment (the "prior" according to Bayesian terminology) and combines it with the few available data. The difficulty lies in translating the expert judgment into the language of quantitative analysis.

We shall give a simple methodology to transform expert judgment in a way that can be processed by Bayesian theory. In any case, we have at our disposal historical data on n policyholders belonging to the same "risk class," that is, with sufficiently similar characteristics so that we may assume that the distribution of their risks is similar. For each policyholder i, we observe a response variable Y_i (e.g., the number of different categories of claims, their severity, or a combination of both). We then make the assumption that this variable is distributed according to a reasonable family of distributions. For illustrative purposes, focusing on frequency, we can consider that the Y_is are distributed according to a Poisson distribution, with unknown parameter θ_0.

The Bayesian approach consists in considering that θ_0 is a hidden factor, that is, a random variable with distribution π that has been generated prior to obtaining the observation and that is not observed directly (but is governing the result of the random experiment that produced the data). The practical meaning of this distribution π has to be related to the knowledge that one has independently of observing the historical data. Typically, the definition of π helps to take into account the fact that an expert gave us a reasonable value for the parameter (here the frequency) that we wish to estimate. Assume that this expert considered that the value of the frequency of our risk class should be around the value θ^*. Then a first possibility is to define the density of the prior distribution as $\pi(\theta)d\theta = c\exp(-\lambda(\theta - \theta^*)^2)d\theta$ (restricted to the set of positive values of θ), where c is a normalizing constant that we will have no need to determine explicitly.

This form of the prior density (and one can think of many others) is designed simply to say that it is very unlikely that the true value is far from θ^*, since the value of the density decreases fast when going away from θ^*. Going back to our assumption that the variables Z_i are Poisson distributed with parameter θ_0, the Bayesian estimator of the frequency will be obtained as the maximizer of

$$L\left(Y_1,\ldots,Y_n,\theta\right) - \lambda\left(\theta - \theta^*\right)^2,$$

where $L(Y_1, \ldots, Y_n, \theta)$ denotes the classical Poisson log-likelihood.[2] In other words, the final estimator will be a penalized maximum likelihood estimator, where the penalization comes from the prior and enforcing the final result to be closer to θ^*. If the expert is accurate, this procedure should reduce the volatility of the estimation versus a situation where we would only use historical data of small size.

Of course, experts may be wrong, and statistical testing procedures can be used to check the consistency of their analysis with data (e.g., Evans and Moshonov 2006). Let us note however that when the size of the information grows – that is, when n tends to infinity – this procedure is designed to give less and less weight to the initial expert judgment. A second issue concerns the choice of the parameter λ. Somehow, this parameter drives the confidence one has in the expert's opinion. Typically, it can be calibrated so that the distribution π is concentrated on an interval $[\theta^* - a; \theta^* + b]$, with sufficiently high probability. Here, a and b correspond to the margin of error given by the expert concerning their analysis.

Although we have explained this approach in detail for the frequency, it can be extended easily to the case of the severity or any other quantity, even if its distribution has several parameters (in which case the distribution of the prior is multivariate). The particular choice of the prior that we made allows one to have a simple criterion to maximize, when it comes to estimating the parameter, but it has no closed-form formula in general. If one wishes to obtain closed-form expressions, specific priors have to be chosen that depend on the family of distributions one wishes to fit on Y_i (see the concept of conjugate priors in Diaconis and Ylvisaker [1979], for instance).

[2] If one considers another distribution – for example, when studying the severity – this formula remains true, but with the log-likelihood of the corresponding family of distributions.

We stress that the approach that we develop below can be used as an expert judgment in the Bayesian vision. A prior distribution can be computed from it and can then be combined with historical data of the final user.

4. A Synthetic Database

We describe here the methodology used to generate the database of cyber claims and how to adapt it to generate portfolios with different structures. We describe how the structure of the portfolio has to be specified in Section 4.1. We also define in Section 4.3 a classification of the potential type of events. For each policyholder in the portfolio, we then simulate the number for each type of claim experienced during the period of interest (typically, one year). We describe the model and how its parameters are specified in Section 4.2. For each claim, the severity is then simulated according to a procedure described in Section 4.4.

4.1. Exposure

The portfolio is composed of policyholders described by a set of characteristics summarized in Table 1.

We intentionally restrict characteristics to a small number. Many other risk factors have impacts on the frequency and severity of cyber events, but the more parameters, the more values to calibrate according to data, as well as difficulties in capturing the interaction between these factors. For example, turnover is likely to have an effect on the probability of a company being targeted. But we choose not to include it in the model because of a lack of statistics linking turnover to attacks and also because it is correlated with the size, sector of activity, and country, potentially repeating some information contained in those variables. On the other hand, we will use the turnover rate to estimate the financial loss associated with cyber events – see Section 4.4.2.

The effects we kept in Table 1 are the most obvious in the context of cyber risk. Let us discuss the motivation for taking them into account:

- **Size:** The size of the company is, first, a factor of visibility. It can attract the attention of hackers. But more than that, size gives information about the attack surface, since the human factor is key in allowing hackers into the systems.

Table 1. Characteristics of policyholders

Fixed Effects	Number of Modalities
Size	4 categories (1–9 employees, 10–49, 49–249, 250+)
Cyber maturity	3 categories (cyber novice, intermediate, expert)
Random Effects	Number of Modalities
Country	8 countries (corresponding to the countries in the scope of Hiscox report)
Sector of activity	14 categories (see the list below)

- **Cyber maturity:** This variable is related to the "cyber hygiene" of the company: how well it is organized in terms of securing its information systems and whether its employees are trained to prevent or respond to cyber incidents. There are several metrics with which to define cyber maturity, also related to the proportion of the budget spent on cybersecurity. We picked categories that are consistent with the definition proposed by the Hiscox report due to the ease of calibrating models according to that reference. However, other metrics can be used as long as one has enough information on their impact on cyber risk.
- **Country and sector of activity:** These variables add information on the degree of preparedness of the victims and on the dependence of the company on its information systems.

4.2. Frequency

The number of claims N_i experienced by the ith policyholder is modeled using a Poisson random variable. The mean parameter is $\lambda_i = -\log(1 - p_i)$, where p_i is the probability of being a victim of at least one event. The choice of a Poisson distribution is a way to materialize the fact that a policyholder can be struck more than once, while the quantity p_i is easiest to calibrate from available data.

First, let us recall that we introduced two random effects (country and sector) that apply to all policyholders of the same category (with respect to these two characteristics). These effects are materialized by two random vectors: Δ_c for the country (its dimension d_c corresponds to the number of countries) and Δ_s for the sector (its dimension d_s corresponds to the number of sectors). The random vector (Δ_c, Δ_s) of size $d_c + d_s$ is simulated using a multivariate centered Gaussian distribution with covariance matrix Σ. The matrix Σ used in our simulation setting is diagonal, with variance 0.0284 for the country and 0.10145 for the sector.

The value p_i is taken as

$$logit(p_i) = \alpha_i + \beta_0 + \beta^T \mathbf{X}_i + \Delta_{c,i} + \Delta_{s,i},$$

where $\Delta_{c,i}$ (resp. $\Delta_{s,i}$) is the coordinate of Δ_c (resp. Δ_s) corresponding to the country (resp. the sector) of policyholder i, and $\mathbf{X}_i$ is the vector of covariates corresponding to the four characteristics of policyholder i.

We introduce a hierarchical relationship between these characteristics in the sense that the fixed effect coefficients (corresponding to size and cyber maturity) depend on the country and the sector within the country. The list of countries is limited to those for which data are available in the Hiscox report (Table 2). The list of sectors was created in accordance with the Organisation for Economic Co-operation and Development categories in order to retrieve data about the proportion and characteristics in the economy of a given country.

Table 2. List of the countries and sectors of activity following the categories present in the Hiscox report

Countries	Sectors of Activity
Belgium France Germany Ireland Netherlands Spain United Kingdom United States	Business services Construction Energy Financial services Food and drink Government and nonprofit Manufacturing Pharma and healthcare Professional services Property Retail and wholesale Information and communication Transport and distribution Travel and leisure

4.3. Types of Events

Once we have defined the number of cyber events an entity has encountered, we determine the type of cyber incidents that struck the policyholder. This is the first step toward establishing a link between an event and its severity, because the consequences of a cyber event are different depending on the type of attack.

To define a cyber event, we have chosen to characterize them by type of attack, denoted by q, and one (or more) entry point(s), denoted by v, except for DDoS attacks, for which we do not specify the entry point. The list of attack types and entry points is given in Table 3.

Here again, the choice of a relatively small number of situations is guided by the need to gather enough data to correctly estimate the probability of occurrence of each type of attack. Obviously, the type of attack will have an impact on the consequences of the attack. Moreover, the type of attack must be distinguished from the "entry point," which is also key information, since it enables an entity to improve prevention and anticipate effects.

Table 3. List of types of attack and possible entry points

Type of Attack q	Entry Point v
DDoS	Phishing email
Ransomware	Credential theft
Loss of data (other than ransomware)	Third party
Business email compromise	Unpatched server
Other	Brute force server credential

In our approach we treat the DDoS attack differently from the others, in the sense that we consider that this specific type of attack has no entry point. DDoS attacks are most often carried out by bots that have not entered the target's system but that generate a saturation of requests that bring down the servers (Thing et al. 2007). When it comes to studying the impact of the attack, we have also developed a specific way of dealing with DDoS attacks, since the interruption of service is supposed to be essentially limited to the time before the bots stop attacking. However, it should be noted that in some cases DDoS attacks are used to distract a victim while another type of attack (e.g., ransomware) is being carried out. We do not take into account this sort of combination of attacks, considering, for simplicity's sake, that the "ransomware" label would be assigned to the event in such a case.

The distribution of the different types of attack we are considering, based on Hiscox statistics, is given in the appendix (Table 18). Note that there may be multiple entry points, as shown in Table 20, where we assign a probability to each type of entry. With regard to this aspect, we adopt a specific treatment for phishing emails. Phishing can essentially lead to the downloading of malicious applications or the theft of credentials. Consequently, in the database we generate, we force phishing to be associated with the theft of identification data in 51% of cases (in line with US statistics).

4.4. Severity

For each incident, we decompose the severity into several components. First, we consider the victim's ability to defend itself from the attack. From an insurance perspective, these situations correspond to "false alarms" since they do not lead to a financial loss. Nevertheless, it is important to monitor them, since they can allow us to project what would happen if hackers improved their rate of success or if, conversely, policyholders increased their level of security and managed to defeat a higher proportion of attacks – see Section 5.3. The probability that the victim can successfully defend itself is given in the appendix (Table 20) and depends on the entity's cyber maturity.

If the attack is successful, to determine the severity, we then distinguish between the volume of exposed data, in the case of data theft or compromise, and the time of business interruption. As already mentioned, a distinction is made between the business interruption caused by DDoS (shorter) and other types of attacks like ransomware.

4.4.1. Size of the Data Breach

The magnitude of a data breach is often expressed in terms of "number of records." To generate the database, we simulate this number R according to the model fitted by Farkas et al. (2021) on the PRC database mentioned in Section 2.3. As mentioned in Section 2.4, the distribution of R has been identified as heavy tailed by several authors, whereas, focusing on the center of the distribution more than on the tail, Eling and Loperfido (2017) considered that a lognormal distribution was adapted.

Therefore, the idea is to distinguish between standard data breaches, which happen with probability $p = 0.84$ according to Farkas et al.'s (2021) model, modeled using a lognormal distribution (i.e., log R is distributed according to a Gaussian distribution with

mean μ and variance σ^2), whose parameters are discussed below, and a large database (probability $1 - p = 0.16$), simulated using a shifted generalized Pareto distribution whose parameters depend on the characteristics.

To transfer the model proposed by Farkas et al. (2021) into our framework, some modifications have been made to take into account the fact that the variables present in the PRC database are not the same as those we want to focus on in the insurance framework. The values of the parameters for the distribution of standard breaches are given in Figure 1. Values and parameters for extreme breaches are shown in Figure 2.

4.4.2. Cost of a Cyber Event

Again, we choose to model the physical consequences of a cyberattack, rather than the attack's true cost, due to the greater amount of available information on the former. In this section, we list proxies that allow for the transformation of those consequences into an approximation of the cost of the event.

The question of the cost of a data breach is relatively well documented. A yearly report produced by IBM Security monitors the evolution of losses related to such incidents – see https://www.ibm.com/reports/data-breach. The 2023 report regroups the financial consequences of a data breach into four categories:

- Detection and escalation (36%)
- Lost business (29%)

Figure 1. Distribution of standard breaches, values of μ and σ for the lognormal distribution (adaptation of Farkas et al. [2021]).

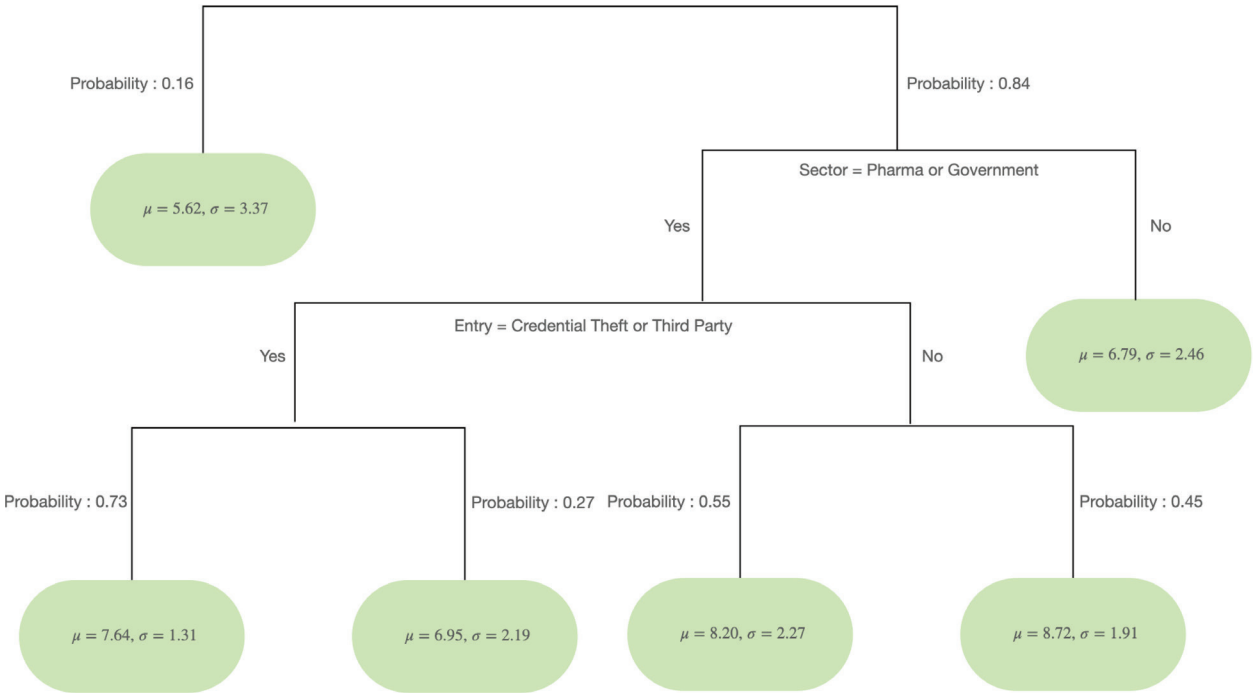

Figure 2. Distribution of extreme breaches, as a mixture of generalized Pareto distributions, which means that $\mathbb{P}(R - u \geq t) = \dfrac{1}{\left(1+\dfrac{\gamma t}{\sigma}\right)^{1/\gamma}}$. The values of γ and σ for the generalized Pareto distribution (adaptation of Farkas et al. [2021]). The threshold is $u = 27999$.

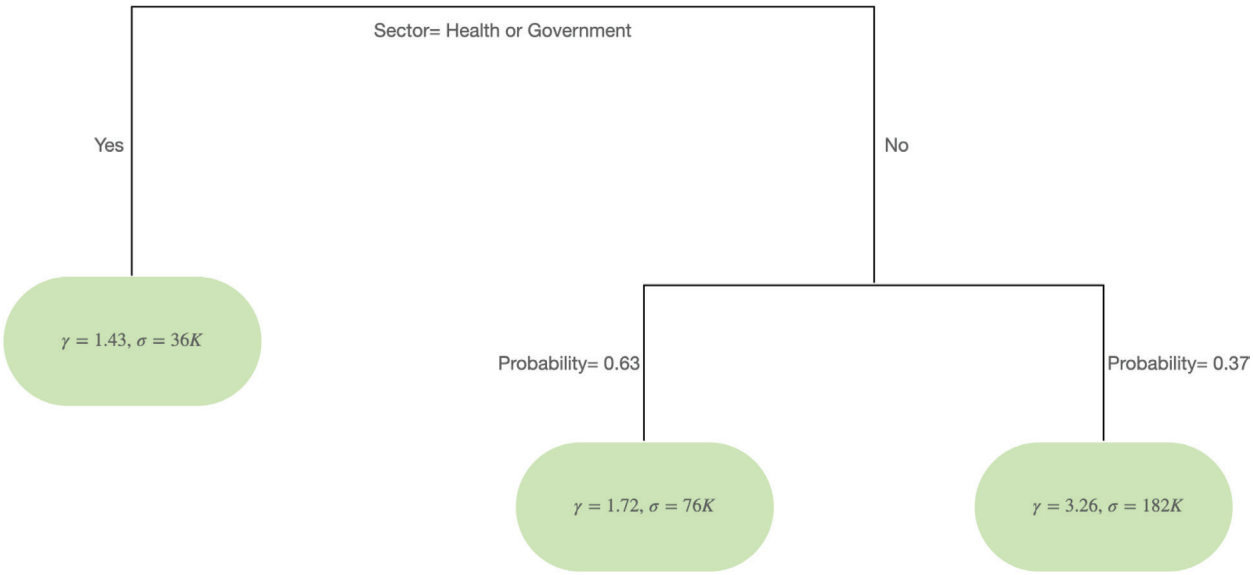

- Post-breach response (27%)
- Notification cost (8%)

However, these categories are relatively imprecise, and the actual impact of the consequences (and the ratios between these different cost segments) should vary considerably depending on the characteristics of the victim.

It is therefore difficult to model the cost of a data breach directly, given the lack of precise data on the subject. Yet, statistical information on the volume of data breaches is made public by various sources, such as the PRC database (see https://privacyrights.org/data-breaches) or the VERIS database (see https://verisframework.org/). Hence, it is easier to model the "number of records," i.e., the volume of data exposed by the attack. Several evaluation formulas have been proposed to transform this information into an economic loss.

If R denotes the number of records and L the financial loss, Jacobs (2014) proposed the model

$$\log L = 7.68 + 0.76 \log R,$$

this formula being calibrated from data from the Ponemon institute. Farkas et al. (2021) proposed an update of that formula that takes into account "mega-breaches" that happened more recently, leading to

$$\log L = 9.59 + 0.57 \log R.$$

Romanosky et al. (2019) proposed a more detailed formula that takes into account the revenue of the victim, the existence of lawsuits after the breach, and several other variables that make the formula more precise but less easy to use and calibrate.

However, the number of records an attack exposes is not the only indicator of its severity. Some DDoS attacks, for instance, do not even involve breaches but still result in damage. Whereas the issue of data exposure has received a lot of attention because of its implications for data confidentiality, business interruption can generate considerable losses. To illustrate the issue of business interruption, we have chosen to model its various stages in detail, as described in Section 4.4.4. We address the specific case of DDoS attacks in Section 4.4.5.

For non-DDoS attacks, an approximation for estimating the loss corresponding to a business interruption lasting T is to compute $turnover \times T/365$, where the victim's turnover (understood as volume of sales) is taken into account. This rough approximation is based on the assumption that the company's turnover is produced uniformly over the year, which is not necessarily true. Furthermore, it does not take into account the support or repair costs generated by the incident. The Lloyd's report *Cloud Down* (https://www.lloyds.com/clouddown) used a similar approximation to assess losses, limiting turnover to that generated by the victim's online activities. Although it may be interesting to use information on the insured's online surface, it should not be forgotten that cyberattacks can also contaminate offline activities, such as manufacturing that is interrupted due to the unavailability of digital tools.

DDoS attacks are different. The interruption is much shorter, usually less than a day. Moreover, they are generally designed to hit the victim at the worst possible moment. Consequently, applying the same formula to deduce the loss would lead to an underestimation of their impact. The Ponemon Institute estimated (in 2012) the average cost of one minute's downtime for a DDoS attack at $22,000, with a wide range (depending on the case, this can be anything from $1 to more than $100,000 per minute). NSFOCUS's *2022 Global DDoS Attack Landscape Report* (https://nsfocusglobal.com/nsfocus-releases-2022-global-DDoS-attack-landscape-report/) provides a detailed overview of the consequences of DDoS attacks.

4.4.3. Duration of a Business Interruption

The overall methodology we use is summarized in Figure 3 and developed in Sections 4.4.4 and 4.4.5.

4.4.4. Non-DDoS Attacks

In the case of non-DDoS attacks, we decompose the business interruption into the following steps:

1. First interruption (stage where the disruption is the highest).
2. Start of a backup plan that reduces (but does not suppress) the damages caused by the attack. In terms of loss, this backup plan can be materialized by a coefficient reducing the cost during this period.

Figure 3. Overall process

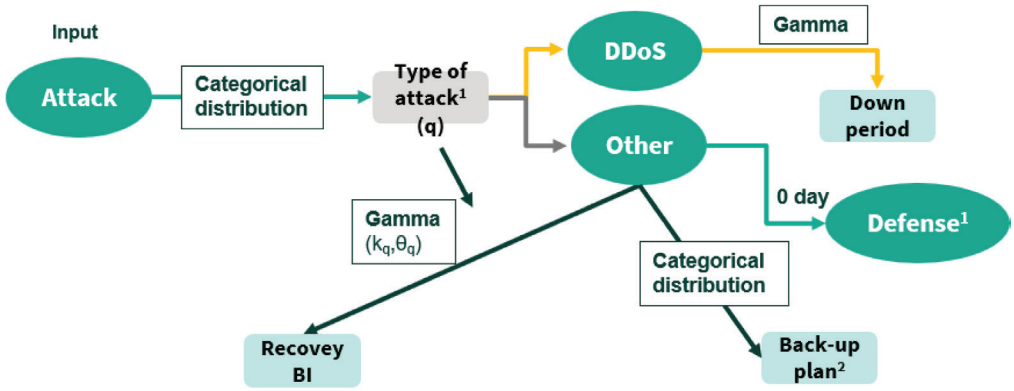

3. Restoration of the information system: the attack has stopped, but the server may not have been restored.
4. Return to normal.

The probability distributions governing these different steps are listed in the appendix (Tables 21 to 23).

4.4.5. Duration of DDoS Attacks

To simulate the duration of DDoS attacks, we referred to the Netcloud 2022 DDoS Threat Intelligence Report (see https://www.netscout.com/threatreport/), based on NETSCOUT's expertise in the field of DDoS. As with many other statistics we use in the paper, these statistics may contain bias as there is no clear traceability of the population that was observed to determine these figures. However, the intention is to create a reasonable benchmark, which needs to be supported by real portfolio data for industrial applications.

The distribution of the duration of the attack is governed by the probabilities in Table 4. Because the table gives only a range of duration (e.g., 10–60 minutes), a uniform distribution is used over the time interval corresponding to the event. It should also be noted that no relevant statistics are available for DDoS attacks of very long duration, corresponding to 12 hours or more of business interruption. Due to the rarity of such events, assumptions must be made to deal with this particular case.

4.5. Pricing: A Detailed Example

Through a Monte Carlo analysis, the synthetic database developed herein can help project the result of a given insurance portfolio. It is also a way to verify the behavior of statistical methods that may be used on real portfolio data.

Table 4. Probability that the different points of entry are involved

Duration of the attack	< 5 min	5–10 min	10–60 min	1–12 h	> 12h
Probability	25%	36%	30%	7%	2%

The parameters we have considered can be used to determine a premium. We describe here, for a given example, how to establish the pure premium, as a function of the chosen parameters. We then explain the main steps involved in moving from this pure premium to the final premium paid by the policyholder. The characteristics of this policyholder are given in Table 5.

Using a classical frequency-severity approach, the pure premium is the product of the expected number of claims and the average severity of a claim. Regarding the frequency of incidents, we obtain a value of $p_i = 0.59$ for the probability of at least one incident and an intensity parameter of $\lambda_i = 0.89$. Recall that, among these incidents, a significant number will not trigger any insurance compensation because the policyholder will succeed in defending itself in certain cases or because the amount of the loss will be lower than the deductible.

To avoid any redundancy, we focus on an insurance contract covering only business interruption and excluding DDoS attacks. If an additional guarantee is linked to the volume of affected data, the pricing techniques are easily extended, following the same path, but based on the model in Section 4.4.1. Here we consider a contract with a US$100,000 deductible.

From Table 19 in the appendix, we see that 29.8% (at most) of attacks are covered, given that we are dealing only with the categories *ransomware* and *other*. Since the policyholder is of *intermediate* cyber maturity, the probability of it defending itself is 9.5%, which leaves only 45.6% of cases where insurance compensation is triggered (provided that it exceeds the deductible). In the case of a claim, if T is the downtime period in days, U is the time at which the backup plan is triggered, and V is the time after restoration of service when things go back to normal, the simplified framework proposed by the Lloyd's report implies that the cost C for the policyholder is given by

$$C = \left(\frac{200}{365}\right) \times \left(\sum_{i=1}^{T} 1_{\{i \le U\}} + 0.40 \times 1_{\{i > U\}} + \sum_{i=T+1}^{T+V} \left(1 - \frac{(i-T-1)}{V-T-1}\right) \times \left(1_{\{T \le U\}} + 0.40 \times 1_{\{T > U\}}\right) \right),$$

where $1_{\{A\}}$ is the indicator function (equal to 1 if the event A appearing within the brackets is realized and 0 otherwise). The ratio 200/365 corresponds to the expected loss of one day of business interruption, following the simplification made by the Lloyd's report mentioned above. After applying the deductible and the parameters mentioned in the appendix,

Table 5. Policyholder used for the illustration

Size	250+ employees
Cyber maturity	Intermediate
Country	US
Sector of activity	Energy
Turnover	US$200 million

we obtain an average cost of US$120,000, which leads to an average severity of $0.298 \times 0.905 \times 120 =$ US$32.36K. The pure premium then becomes $0.89 \times 32.36 =$ US$28.80K.

Note that here we have considered only a specific type of event. To cover several types of events, the pure premium is obtained by adding together the premiums corresponding to each type of event covered. It should also be noted that the value is highly dependent on how downtime is converted into a financial loss. The approach we adopt is approximate, and given the heterogeneity of the situation of policyholders, this part should be studied in greater depth. This is particularly the case for DDoS attacks, which are shorter in terms of business interruption but are associated with a greater loss per unit of time. Another important issue is that we are calculating a pure premium here, which assumes that there is independence between frequency and severity. A loading factor must be applied to account for situations where that assumption may not be true. These instances are described in the next section of this paper.

5. Applications

5.1. Risk Transfer Solutions

Insurance has already been considered in the preceding sections, and we therefore focus here on alternative strategies. Indeed, in response to a rise in the frequency and severity of cyber claims, a number of insurers have reduced their offer or moved it to higher layers. This has resulted in more limited capacity and a hard insurance market, which has increased demand for alternatives. Faure and Nieuwesteeg (2018) and Eling et al. (2021) carefully discuss strategies for cyber risk management including retention, pooling, and transfer.

The first risk management strategy is retention. By designing effective security protocols and ex post risk mitigation processes, firms can keep a substantial part of their cyber risk exposure. This is particularly attractive for dealing with relatively small cyber risks – i.e., those that would result in only limited damage if the risk were to materialize – or for large corporations.

Beyond retention, cyber risk pooling offers promising opportunities. Recall that pooling is in essence risk sharing between participants without transfer to an insurer. It can be implemented through the establishment of a mutual insurer. We will discuss a recent initiative in that direction. At the end of 2022, a new cyber mutual insurer called MIRIS (Mutual Insurance and Reinsurance for Information Systems) was licensed by the Belgian regulator. Its purpose was to underwrite direct cyber insurance on behalf of its owner-members based in the European Union and European Economic Area.

In addition to cyber risk coverage, MIRIS aims to promote cyber risk protection. Within a cyber risk pool, there are indeed strong incentives to share current knowledge and best practices, especially among members facing similar exposure to cyberattack. In that respect, it turns out that MIRIS's founding members include the German chemical group BASF and its Belgian peer Solvay, together with 10 other large industrial corporations.

If claims exceed deposits and exhaust the available risk capital, participants will be subject to "supplementary calls," similar to the mechanism used by P&I (property and casualty) clubs. Because the performance of MIRIS depends on maintaining an aggregate risk profile better than the market, all interests are aligned and pooling favors active collaboration between participants to improve prevention and effective risk management. MIRIS's members' chief information security officers and insurance managers screen new members for both their financial strength and their cyber risk management capability. They also promote cyber risk management best practices between the members to ensure the ongoing high quality of the members' risk management.

Establishing captive insurance structures can also offer an effective solution to meet sizable cyber exposures. Captives are especially helpful in the current hard cyber insurance market, offering capacity at attractive conditions and helping to keep the overall cost of cyber risk management under control.

5.2. Parametric Insurance

Among the possible risk transfer solutions, parametric insurance is often mentioned as promising when it comes to covering complex risks. The principle of parametric insurance consists in not paying the exact amount of loss experienced by the victim: if this loss is L, the parameter P is a quantity that is correlated to L, but that is supposed to have the following advantages:

- P can be easily measured soon after occurrence of a claim, while L may require a detailed expertise before being known exactly.
- From P, one can estimate with good precision L, which means that the compensation may not be perfect for the policyholder, but satisfactory enough.
- One has gathered data on P, so actuarial modeling of insurance products based on P is much more easy.

This last characteristic is particularly appealing for the insurer in a context where data on cyber are scarce. On the other hand, the policyholder may elect to receive a compensation that is not its true loss in exchange for a faster payment. Because the parameter P is available soon after occurrence of the claim, the compensation can be paid much faster. This celerity is important in the context of cyber, where the victim needs funds to quickly restore its information systems so as to return to normal activity as soon as possible.

But of course the key issue is to find an available parameter that meets the policyholder's needs. In addition, this parameter must be sufficiently comprehensible to the potential customer. For the time being, this considerably limits what can be done in terms of designing parametric products in the field of cyber insurance. Typically, existing solutions focus on the duration of the business interruption caused by an attack. While useful, this covers only a specific part of cyber insurance and may not address the whole problem, as the impact of business interruption time, while relevant, may vary depending on the specifics of the insured, even after filtering for standard risk factors such as the victim's turnover or sector

of activity. In addition, coverage for extreme losses can be disappointing for the customer—see Lopez and Thomas (2023) and an example of agricultural insurance in Johnson (2021).

In summary, parametric products are promising for covering cyber risk, as they can meet the needs of both the insurer (better anticipation thanks to simplification of the problem) and the insured (rapid compensation). However, they should probably be seen as tools that need to be complemented by other types of products to cover cyber risks effectively.

5.3. Prevention

Prevention is always crucial in insurance for many reasons, from avoiding the economic burden of claims to maintaining a positive image of the insurer as being committed to helping its customers. But it is especially important in the context of an emerging risk such as cyber, because it is also a way to anticipate the evolution of the risk. The cost of prevention, in cyber, may be relatively high, due to the poor degree of education (although improving) of the population. It can be thought of as a particular form of risk transfer, because budget is spent to help cybersecurity suppress some part of the risk. However, it is hard to evaluate the effect of one dollar spent on prevention in terms of reduction of risk. Moreover, prevention can be implemented in various ways. The most obvious is to invest in IT solutions that better detect intrusions and/or can contain them efficiently. Similarly, the human factor is also key in cybersecurity, as shown in Rahman et al. (2021). Improving the organization's management or developing a culture of cyber risk awareness among employees can considerably contribute to risk reduction.

A good example is the case of phishing. IT solutions, such as the introduction of firewalls and spam filters, can reduce phishing. However, completely suppressing these malicious emails is impossible. Moreover, the development of generative artificial intelligence represents, in this field, a threat, because of its ability to generate more elaborate fraudulent emails that may be harder to detect by automated systems (Teichmann 2023; Neupane et al. 2023). At some point, reduction of this risk requires people changing their behavior. For instance, credential theft usually happens due to an error made by an individual, even in the case of brute force attacks (that is, multiple attempts by bots to crack passwords) that succeed because of insecure passwords.

Using the database we developed, in the following examples we address (partially and approximately) some questions about the costs and optimization of prevention.

Example 1: Effect of DDoS protection. Protection against DDoS attacks is essentially a matter of IT. Better organized information systems, for example, minimize the attack surface and reduce the potential ways by which attackers may strike. Investing in servers with a traffic capacity sufficient to meet the company's needs is, of course, key. Also, tools to detect abnormal requests sent to servers can help ensure a fast reaction, and firewalls may be implemented to stop attacks.

We consider here the case where a company increases its level of protection against DDoS attacks by a certain percentage. This translates into a change in the probability of defeating the attack when the event is of the DDoS type. In Table 6, we show the impact on the average

Table 6. Impact of improving defense against DDoS attacks. The last three columns refer to the impact of defeating x% more cyberattacks. The numbers are expressed in percentage of the average cyber loss (all causes) of the victim.

Maturity	10% Reduction	15% Reduction	25% Reduction
Novice	2.44%	2.55%	2.78%
Intermediate	3.09%	3.23%	3.52%
Expert	3.58%	3.74%	4.07%

loss suffered by the company generated by cyber events (of all types). Note that in this example (and similarly in Tables 6 to 10), we only consider the average loss. But the process of simulating a pseudo-database that we considered earlier can help us to obtain more information about the effect on the loss distribution (variance or quantiles, for example).

Example 2: Prevention of phishing attacks. We distinguish between two cases to measure the impact of prevention in this area: phishing attacks that also involve credential theft and those that do not. In the first instance, a company can reduce the number of attacks by using better firewalls and spam detectors. Again, we express improvement as percentage increase in successful defenses against phishing attacks (Table 7). Involvement of credential theft means that some workers in the company were mislead by the fraudulent email and gave their credentials to the hackers. Education of workers is a way to reduce the impact of such events. On the other hand, other types of phishing attacks, such as the downloading of malicious elements, can be overcome in part by education and in part by technical solutions that detect malicious software and/or prevent employees from installing malicious elements without the system administrator's consent (Table 8).

Example 3: Efforts in data security and crisis management. In the previous examples, we considered the improvements generated by reducing the number of cyberattacks. But another means of prevention is to better prepare for the consequences of such attacks. We consider two types of action:

- creation of a backup plan (Table 9); and
- an x% reduction in the time before going back to normal (Tables 10 and 11).

Table 7. Impact of improving defense against phishing attacks with credential theft (second number corresponds to attacks that use credential theft and phishing as sole entry points). The last three columns refer to the impact of defeating x% more cyberattacks. The numbers are expressed in percentage of the average cyber loss (all causes) of the victim.

Maturity	10% Reduction	15% Reduction	25% Reduction
Novice	2.65%/0.92%	2.77%/0.97%	3.01%/1.05%
Intermediate	3.36%/1.18%	3.51%/1.23%	3.82%/1.34%
Expert	3.89%/1.36%	4.06%/1.42%	4.41%/1.55%

Table 8. Impact of improving defense against phishing attacks that do not involve credential theft (second number corresponds to the attacks that use as sole entry point). The last three columns refer to the impact of defeating x% more cyberattacks. The numbers are expressed in percentage of the average cyber loss (all causes) of the victim.

Maturity	10% Reduction	15% Reduction	25% Reduction
Novice	2.55%/0.89%	2.66%/0.93%	2.89%/1.01%
Intermediate	3.23%/1.13%	3.37%/1.18%	3.67%/1.28%
Expert	3.23%/1.31%	3.37%/1.37%	3.67%/1.49%

Table 9. Reduction of loss generated by business interruption if a backup plan is used. We consider the cases of 1 day, 3 days, and 5 days during which the technical solution to fix the vulnerability has not been found; full restoration of the activity can continue after this initial period.

Turnover	1 Day	3 Days	5 Days
≥ $1 billion	29%	31%	32%
< $1 billion	24%	26%	27%

Table 10. Reduction of the loss generated by business interruption if a backup plan is used. We consider the cases of 1 day, 3 days, and 5 days during which the technical solution to fix the vulnerability has not been found; full restoration of the activity can continue after this initial period.

Reduction of the Time of Recovery	1 Day	3 Days	5 Days
50%	32%	18%	13%
25%	16%	9%	6.5%

Table 11. Reduction of the loss generated by business interruption in the case of a 50% reduction of the average amount of time before going back to normal, in presence of a backup plan

Turnover	1 Day	3 Days	5 Days
≥ $1 billion	82%	60%	40%
< $1 billion	81%	59%	39%

5.4. Accumulation Stress Scenarios

We consider two types of scenarios: (i) a contagious cyberattack, and (ii) the failure of a digital solution (e.g., cloud) shared by a significant part of the portfolio.

5.4.1. Contagious Cyberattacks

The pattern of a contagious cyberattack is very similar to a biological epidemic, with different kinetics. Among the famous such cases, WannaCry (Mohurle and Patil 2017) and NotPetya (Fayi 2018) probably received the most media attention, but some, like the Love Bug (Cohen and Anderson 2000), occurred in the early days of the internet. Hillairet et al. (2022) proposed a way to model the impact of such episodes (see also Hillairet and Lopez 2021) on a cyber insurance portfolio based on the modeling of three aspects:

1. The spread of the cyber epidemic, based on compartmental models classically used in epidemiology
2. The time of immediate assistance required by the victim (different from the total time required to fully recover from the attack)
3. The time required to identify a way to correct the vulnerability used by the attack, and how fast it is adopted by policyholders

We mostly focus here on the first aspect, but we explain how the other two can enrich this scenario of the impact of a contagious cyberattack.

Regarding the spread of the cyber epidemic, a full description of the compartmental model is provided in Section A.1. To formulate a scenario of attack, one first needs to model the behavior of the attackers: is the initial attack targeted on a specific (or a few) class(es) of policyholders, or does it strike uniformly the whole exposed population? Next, to take into account the propagation (potentially to classes of policyholders that were not the primary target of the attackers), we need to introduce the network structure between the categories of policyholders. That network structure is clearly hard to come by. Understanding precisely the connections between all the actors involved in the propagation is beyond our reach. Hillairet et al. (2022) propose that we consider the network structure only at a macroscopic level, that is, with average values of connections between categories of victims, and then introduce some randomness to take into account the approximation made at this stage. Finally, we take into account the strength of the contagion with, essentially, two parameters (fully described in Section A.1), one of them (denoted by β) giving the propensity of an infected system to propagate the virus and the other (denoted by γ) materializing how fast the propagation can be contained inside a given system and stopped from spreading to other victims.

From these inputs, one can rapidly estimate the number of infected policyholders within the portfolio in each category of policyholder.

Table 12 shows results from Hillairet et al. (2022) of a simulation of a cyberattack of the same intensity as WannaCry using this type of contagion model. The output of the scenario

Table 12. Proportion in each sector of companies affected by the epidemic, depending on the targeted sector

Targeted Sector	Mining	Manufacturing	Energy	Construction	Services
Uniform attack	1.06%	4.11%	0.99%	2.07%	8.86%
Attack on mining	99.70%	12.69%	1.36%	5.49%	20.37%
Attack on manufacturing	1.02%	16.01%	0.66%	3.05%	16.58%
Attack on energy	0.93%	5.96%	64.08%	2.35%	12.93%
Attack on construction	0.33%	2.49%	0.21%	6.60%	5.72%
Attack on services	0.25%	2.59%	0.21%	1.01%	7.84%

differs depending on the sector of the economy struck by the attackers (here, to simplify, only five industrial sectors are considered); the contagion then spreads to other sectors of activity. We refer the reader to Hillairet et al. (2022) for a full discussion of these scenarios, their calibration, and the possibility of reducing their impact depending on the quality of the prevention.

5.4.2. Cloud-Type Attack

The cloud outage scenario is often mentioned as one of the more concerning ones for the insurance sector, and it is at the core of many stress scenarios (see EIOPA 2023). The report *Cloud Down: Impacts on the US Economy*, co-produced by Lloyd's and AIR (see https://www.lloyds.com/clouddown), lists some significant cloud events occurring between 2012 and 2017. None of the incidents covered in this 2018 report was malicious, but their impacts demonstrate the severity of the consequences of such an event. The CloudNordic incident (August 2023) is a recent example of a ransomware attack targeting a cloud provider. All of the infected servers were wiped by the hackers, leading to a permanent loss of data for the customers involved. For now, the exact losses and responsibilities are not publicly known; the impact on CloudNordic's customers may be mitigated by a diversification in their data storage solutions, with potential replications of critical information that were saved in other structures. But, from the insurance perspective, the cloud attack scenario is concerning because of the attack's ability to instantly strike a large number of policyholders.

The first recommendation for anticipating the impact of such an event on a portfolio is to determine the exposure and how diversified the portfolio is. A first step can be to retrieve information about the cloud solution(s) used by the policyholders. In Table 13, we reproduce numbers from Flexera's *2023 State of the Cloud Report* (Flexera 2023) showing the average distribution of public cloud solutions between companies (with, of course, some companies relying simultaneously on multiple cloud solutions). It is important to distinguish between public and private clouds. Public clouds rely on infrastructure shared by a large number of users – these are the most concerning when it comes to triggering "systemic" events. Private clouds, on the other hand, suppose a single data center for any single customer.

Table 13. Adoption rates of different cloud providers (restricted to cases where these solutions are used to run significant workloads). Source: Flexera (2023).

Cloud Provider	Adoption Rate
Azure	41%
AWS	47%
Google Cloud Platform	17%
Oracle Cloud Infrastructure	8%
IBM Cloud	8%
Alibaba Cloud	5%
Other	5%

The risk of a systemic event striking a private cloud is not completely absent. If hackers discover a vulnerability in the software of a given cloud provider, customers using that cloud solution become more likely to be attacked soon after this discovery. However, in the case of a public cloud, a single attack instantly generates a large number of victims. In the private cloud case, the efforts of the hackers are more important, since they must elaborate a distinct attack for each of their targets. In this sense, this case is close to the situation considered in Section 5.4.3.

The precise consequences of a cloud outage on a given policyholder are hard to determine, as they depend strongly on the specifics of each case. Because a precise analysis of the consequences on each and every policyholder appears impossible, one needs to determine a distribution for the time of interruption and then link that interruption time with an economic loss.

More precisely, and following a methodology developed in the aforementioned *Cloud Down* from Lloyd's and AIR (https://www.lloyds.com/clouddown) to assess the potential impact of a (public) cloud attack scenario on a given provider, one can do the following:

1. Determine the proportion of the portfolio affected by the loss. For example, assuming that the distribution of cloud solutions in our portfolio is similar to the average distribution in the global population described by Table 13, then an attack against AWS would lead to 47% of policyholders affected. However, we should note that one might regard such an estimate as pessimistic, since, for the biggest cloud providers, it seems unlikely that all of their servers would go down simultaneously.
2. Determine the duration of the outage.
3. Determine the time needed for the victim to adapt to the outage – that is, its ability to deploy a backup plan. Depending on the preparedness of the victim, installing this temporary solution may not take the same amount of time. An efficient backup plan, implemented soon after the business interruption, should reduce the final impact.

Table 14. Example of losses from a cloud interruption of a company generating US$500 million turnover, according to parameters from Lloyd's and AIR's report *Cloud Down* (see Section A.2.3). After the start of the backup plan, the loss is assumed to be reduced by 40%.

Initial Interruption	Start of Backup Plan	Date for Return to Normal	Loss
1 day	None	Day 3	$4.1M
3 days	Day 2	Day 5	$4.7M
5 days	Day 3	Day 7	$6.8M

4. Determine how long it will take to go back to normal after the resolution of the outage.
5. Determine the typical cost of a day of business interruption depending on the nature of the entity and whether a backup plan has or has not been implemented.

Table 14 illustrates such an approach, taking as input distributions those calibrated by the Lloyd's and AIR report for steps 2 to 4. The precise list of these assumptions is given in Section A.2.3. Regarding step 5, the Lloyd's and AIR report proposes a rough quantification of the loss. According to their formula, the loss is proportional to a company's turnover and the percentage of e-business undertaken by the company. Although we use this assumption to obtain the losses in our illustration in Table 14, this ignores the fact that business interruption caused by a cloud outage does not only interrupt online activities – during the WannaCry episode, car factories for Renault were paralyzed by the cyberattack even though their assembly was not a matter of e-business. Clearly, a better understanding of the consequences of cloud outages has to be developed to improve the accuracy of such scenarios.

5.4.3. Blind Spots

Unidentified shared vulnerabilities are blind spots in the context of cyber risk. The idea is similar to the cloud attack: a weakness in a software solution shared by a large number of policyholders. This can lead to either an increased number of attacks by hackers trying to exploit the vulnerability or even a systemic event. Compared with a cloud breakdown, designing such scenarios is much more delicate. The insurer has no idea about the exposure of its portfolio: it is impossible to list all of the software used by the policyholders and identify which one(s) may represent a danger. Moreover, it is hard to determine how critical each kind of software is. These situations are difficult to take into account in scenarios but should be kept in mind as additional vulnerabilities.

6. Conclusion

From an insurance market perspective, cyber risk is difficult to quantify because of the novelty of this line of business and because of the rapid evolution of such risk in a context where too little information is available to provide a basis for risk analysis. In this report, which can be seen as an extension of the work of Bean (2020), we discussed the challenges and difficulties in the evaluation, transference, and management of cyber risk. To reduce

the gap caused by lack of data, we built a database of fictitious events that one can use to benchmark actuarial techniques and as a preliminary expert judgment to be combined with real data. We created the database using available public data, with the aim of providing the following:

- Content that is driven by a practical understanding of the nature of cyber risk and the desire to model specific important aspects of cyber risk (ransomware, DDoS, business interruption, etc.).
- An open and reproducible methodology. Because the quality of the data used for this particular analysis can undoubtedly be improved, it is important that the approach be flexible enough to generate a database similar in structure but consistent with the point of view and perspective of the final user.
- A methodology that can be easily updated. As a corollary to the previous point, and given the fast evolution and ever-changing landscape of cyber risk, we do not expect the parameters used in this database to be accurate in a few years. Periodic updating will be needed. Our methodology is designed to ease the updating process.

On the other hand, the events simulated with this database do not capture cyber catastrophes, especially accumulation scenarios. For this reason, we discussed various ways to calibrate such scenarios. Although we expect the methodologies presented herein to be helpful in defining stress scenarios, we emphasize that the lack of experience and historical data regarding cyber catastrophes constitutes a challenge for future work. Although we mentioned spectacular examples of the large impacts that cyber may have, all of these events have hitherto been contained. Even the WannaCry and NotPetya episodes, which generated significant losses, appear of low intensity compared with what one might imagine in a potential "cyber-epidemiology," for instance. Therefore, we would stress, as a conclusion of our work, the need to adapt these techniques to the future evolution of the landscape of cyber risk.

References

AMRAE (Association pour le Management des Risques et des Assurances de l'Entreprise). 2023. *LUCY: Light Upon Cyber Insurance.*

Bean, M. A. 2020. *Exposure Measures for Pricing and Analyzing the Risks in Cyber Insurance.* CAS research paper. Casualty Actuarial Society.

Beirlant, J., Y. Goegebeur, J. Segers, and J. L. Teugels. 2004. *Statistics of Extremes: Theory and Applications.* Wiley.

Cohen, D. R., and R. D. Anderson. 2000. "Insurance Coverage for 'Cyber-Losses.'" *Tort & Insurance Law Journal* 35 (4): 891–927.

Daley, D. J., and J. Gani. 2001. *Epidemic Modelling: An Introduction*. Cambridge Studies in Mathematical Biology, no. 15. Cambridge University Press.

Diaconis, P., and D. Ylvisaker. 1979. "Conjugate Priors for Exponential Families." *Annals of Statistics* 7 (2): 269–81.

Edwards, B., S. Hofmeyr, and S. Forrest. 2016. "Hype and Heavy Tails: A Closer Look at Data Breaches." *Journal of Cybersecurity* 2 (1): 3–14.

EIOPA (European Insurance and Occupational Pensions Authority). 2023. *Methodological Principles of Insurance Stress Testing—Cyber Component.* July 11.

Eling, M., and N. Loperfido. 2017. "Data Breaches: Goodness of Fit, Pricing, and Risk Measurement." *Insurance: Mathematics and Economics* 75:126–36.

Eling, M., M. McShane, and T. Nguyen. 2021. "Cyber Risk Management: History and Future Research Directions." *Risk Management and Insurance Review* 24 (1): 93–125.

Evans, M., and H. Moshonov. 2006. "Checking for Prior-Data Conflict." *Bayesian Analysis* 1 (4): 893–914.

Everson, D., L. Cheng, and Z. Zhang. 2022. "Log4shell: Redefining the Web Attack Surface." In *Proceedings Workshop on Measurements, Attacks, and Defenses for the Web.* https://dx.doi.org/10.14722/madweb.2022.23010.

Farkas, S., O. Lopez, and M. Thomas. 2021. "Cyber Claim Analysis Using Generalized Pareto Regression Trees with Applications to Insurance." *Insurance: Mathematics and Economics* 98:92–105.

Faure, M., and B. Nieuwesteeg. 2018. "The Law and Economics of Cyber Risk Pooling." *New York University Journal of Law Business* 14 (3): 923–63.

Fayi, S. Y. A. 2018. "What Petya/NotPetya Ransomware Is and What Its Remidiations Are." In *Information Technology—New Generations*, 93–100. Springer.

Flexera. 2023. *Flexera 2023 State of the Cloud Report.*

Hagen, S., M. Seibold, and A. Kemper. 2012. "Efficient Verification of IT Change Operations or: How We Could Have Prevented Amazon's Cloud Outage." In *2012 IEEE Network Operations and Management Symposium*, 368–76. IEEE.

Hillairet, C., and O. Lopez. 2021. "Propagation of Cyber Incidents in an Insurance Portfolio: Counting Processes Combined with Compartmental Epidemiological Models." *Scandinavian Actuarial Journal* 2021 (8): 671–94.

Hillairet, C., O. Lopez, L. d'Oultremont, and B. Spoorenberg. 2022. "Cyber-Contagion Model with Network Structure Applied to Insurance." *Insurance: Mathematics and Economics* 107:88–101.

Hiscox. 2022. *Hiscox Cyber Readiness Report 2022.*

Jacobs, J. 2014. "Analyzing Ponemon Cost of Data Breach." Data Driven Security, December 11.

Johnson, L. 2021. "Paying Ex Gratia: Parametric Insurance after Calculative Devices Fail." *Geoforum* 125:120–31.

Lallie, H. S., L. A. Shepherd, J. R. Nurse, A. Erola, G. Epiphaniou, C. Maple, and X. Bellekens. 2021. "Cyber Security in the Age of COVID-19: A Timeline and Analysis of Cyber-Crime and Cyber-Attacks during the Pandemic." *Computers & Security* 105: 102248.

Li, Z., M. Liang, L. O'Brien, and H. Zhang. 2013. "The Cloud's Cloudy Moment: A Systematic Survey of Public Cloud Service Outage." Preprint, arXiv, https://doi.org/10.48550/arXiv.1312.6485.

Lopez, O., and M. Thomas. 2023. "Parametric Insurance for Extreme Risks: The Challenge of Properly Covering Severe Claims." Preprint, arXiv, https://doi.org/10.48550/arXiv.2301.07776.

Maillart, T., and D. Sornette. 2010. "Heavy-Tailed Distribution of Cyber-Risks." *European Physical Journal B* 75 (3): 357–64.

Mikosch, T., and A. V. Nagaev. 1998. "Large Deviations of Heavy-Tailed Sums with Applications in Insurance." *Extremes* 1:81–110.

Mohurle, S., and M. Patil. 2017. "A Brief Study of Wannacry Threat: Ransomware Attack 2017." *International Journal of Advanced Research in Computer Science* 8 (5): 1938–40.

Neupane, S., I. A. Fernandez, S. Mittal, and S. Rahimi. 2023. "Impacts and Risk of Generative AI Technology on Cyber Defense." Preprint, arXiv, https://doi.org/10.48550/arXiv.2306.13033.

Ohlsson, E., and B. Johansson. 2010. *Non-Life Insurance Pricing with Generalized Linear Models*. Springer.

Rahman, T., R. Rohan, D. Pal, and P. Kanthamanon. 2021. "Human Factors in Cybersecurity: A Scoping Review." In *Proceedings of the 12th International Conference on Advances in Information Technology*, 1–11.

Reece, M., T. E. Lander Jr., M. Stoffolano, A. Sampson, J. Dykstra, S. Mittal, and N. Rastogi. 2023. "Systemic Risk and Vulnerability Analysis of Multi-Cloud Environments." Preprint, arXiv, https://arxiv.org/pdf/2306.01862.

Romanosky, S., L. Ablon, A. Kuehn, and T. Jones. 2019. "Content Analysis of Cyber Insurance Policies: How Do Carriers Price Cyber Risk?" *Journal of Cybersecurity* 5 (1). https://doi.org/10.1093/cybsec/tyz002.

Teichmann, F. 2023. "Ransomware Attacks in the Context of Generative Artificial Intelligence – An Experimental Study." *International Cybersecurity Law Review* 4:399–414.

Thielman, S. 2016. "Yahoo Hack: 1bn Accounts Compromised by Biggest Data Breach in History." *The Guardian*, December 15.

Thing, V. L., M. Sloman, and N. Dulay. 2007. "A Survey of Bots Used for Distributed Denial of Service Attacks." In *New Approaches for Security, Privacy, and Trust in Complex Environments: Proceedings of the IFIP TC 11 22nd International Information Security Conference (SEC 2007), 14–16 May 2007, Sandton, South Africa*, 229–40. Springer.

Wan, K. S. 2020. "NotPetya, Not Warfare: Rethinking the Insurance War Exclusion in the Context of International Cyberattacks." *Washington Law Review* 95 (3): 1595–1620.

Wolff, J. 2021. " 'Cyberwar by Almost Any Definition': NotPetya, the Evolution of Insurance War Exclusions, and Their Application to Cyberattacks." *Connecticut Insurance Law Journal* 28 (1): 85–129.

Woods, D. W., and J. Weinkle. 2020. "Insurance Definitions of Cyber War." *The Geneva Papers on Risk and Insurance—Issues and Practice* 45:639–56.

Appendix

A.1. Summary of the Epidemiological Model Used for Stress-Testing

The situation of a policyholder is described by the time T of infection by the cyber virus (potentially infinite). It is common to describe the distribution of a time variable by its hazard rate function, that is

$$\lambda_T(t) = \lim_{dt \to 0} \frac{\mathbb{P}\left(T \in [t, t+dt] \middle| T \geq t\right)}{dt}.$$

Let us note that since T can be infinite for some policyholders that never get infected, we do not necessarily have $\int_0^\infty \lambda_T(t)dt = \infty$ as is the case for standard random variables. To follow approaches developed in behavioral epidemiology (i.e., epidemiology where a small subgroup is extracted from a much larger population in which the epidemic spreads), this hazard rate function is obtained from an epidemiological diffusion model.

The SIR model is one of the more simple epidemiological models used to consider the spread of a contagious disease (Daley and Gani 2001). Transposed to the case of a cyber epidemic, the exposed population (supposed to be much larger that the insurance portfolio) evolves, through time, to three states:

- *S*, for susceptible, that is, not infected but potentially infected in the future;
- *I*, for infected, which in our context means contagious, that is, contributing to the spread of the cyber virus; and
- *R* for recovered, which here would rather mean "removed," in the sense that the victim does not contribute to the spread of the virus anymore but may still not be in good shape and may require assistance for a long time before returning to its previous level of activity.

To take into account the heterogeneity between groups of policyholders, the policyholders are gathered into categories (in the example of Hillairet et al. [2022], these categories correspond to business types). For the category j, the total number of individuals with characteristics j is supposed to be constant, and the number of susceptibles in category j (resp. infected) (resp. removed) at time t is $s_j(t)$ (resp. $i_j(t)$) (resp. $r_j(t)$). The evolution of this system of functions is described by the following system of differential equations:

$$\frac{ds_j(t)}{dt} = -\left\{\alpha_j(t) + \sum_{k=1}^{d} \beta_{k,j} i_k(t)\right\} s_j(t),$$

$$\frac{di_j(t)}{dt} = \left\{\alpha_j(t) + \sum_{k=1}^{d} \beta_{k,j} i_k(t)\right\} s_j(t) - \gamma_j i_j(t),$$

$$\frac{dr_j(t)}{dt} = \gamma_j i_j(t).$$

Here, the functions $\alpha_j(t)$ describe the first burst of attacks, which depends on the behavior of the actors. Typically, the attack can strike a certain category of victim at a certain rate during a first step (relatively short before being stopped), and then the contagion spreads. Among the coefficients of the model, $\beta_{k,j}$ describes the way category k contaminates category j. The coefficients γ_j materialize how fast a victim does not contribute to the infection anymore. Hillairet et al. (2022) discuss in detail how to calibrate these parameters with realistic values to mimic a WannaCry-type incident.

These functions (s_j, i_j, r_j) are totally defined by the parameters mentioned above, but there is no closed formula. However, they can be retrieved through numerical approximation with a good precision. The idea is then to define $\lambda_T(t) = \{\alpha_j(t) + \sum_{k=1}^{d} \beta_{k,j} i_k(t)\}$.

With the hazard rate function of T at hand, it is easy to simulate individual trajectories. It is interesting to do so to anticipate how to act throughout the whole crisis. It would be important to know how many policyholders will be affected at the same time. If there are too many, then it will be impossible to bring them assistance simultaneously, and the costs may increase. However, to reduce computation time, one can also be interested in simply knowing the final number of victims. Hillairet et al. (2022) proposed a simple and fast numerical method that does not rely on directly solving the differential system.

A.2. Values of the Calibrated Parameters Used in the Simulations

A.2.1. Random and Fixed Effects

Table 15. List of fixed effect coefficients. The reference category is size 1–9, maturity cyber expert.

Modality	Coefficient
Reference	–0.23631
Size 10–49	0.36461
Size 50–249	0.50064
Size 250+	1.00998
Cyber novice	–0.79035
Cyber intermediate	–0.28354

Table 16. List of country effect coefficients

Country	Coefficient
Belgium	–0.333734601
France	–0.171602822
Germany	–0.273460766
Ireland	–0.214828664
Netherlands	0.005241495
Spain	–0.144006452
UK	–0.473586330
US	–0.278991765

Table 17. List of country effect versus sector

	Sector												
Country	Accommodation	Construction	Energy	Finance	Health	Information	Manufacturing	Public	Real Estate	Services	Transportation	Travel	Wholesale
Belgium	–0.3452	–0.1895	–0.1966	0.0768	–0.5479	–0.1662	–0.6611	–0.1753	–0.1050	–0.2282	–0.1795	–0.2385	–0.4734
France	–0.4494	–0.6334	–0.1005	–0.1529	–0.2708	0.0345	0.0242	–0.3348	–0.1201	0.1513	–0.4055	–0.2781	–0.2993
Germany	–0.3259	–0.3166	0.2065	0.0077	–0.0494	–0.2179	–0.2108	–0.4253	–0.1981	–0.6293	–0.3273	–0.3367	–0.3855
Ireland	–0.1813	–0.4857	–0.1630	0.0888	–0.5687	–0.3289	–0.0931	–0.3434	–0.2792	0.2288	–0.1435	–0.3974	–0.3266
Netherlands	–0.2148	–0.2998	–0.2173	0.2738	–0.2332	–0.1914	–0.3487	–0.2327	–0.3052	–0.0605	0.1925	–0.2069	–0.3418
Spain	–0.4374	–0.2553	–0.2583	0.1130	–0.2198	–0.1061	–0.1494	–0.0614	–0.2121	–0.1560	–0.2334	–0.1971	–0.5601
UK	–0.1636	0.0382	–0.5831	–0.4259	0.0386	0.2537	–0.0363	–0.3932	–0.6702	–0.5193	–0.3118	–0.6625	–0.5068
US	–0.2032	–0.5031	–0.1059	0.1416	0.0477	–0.0397	–0.1258	–0.4280	–0.3311	–0.7657	–0.2155	–0.1905	–0.5096

A.2.2. Parameters Related to the Type of Attack

Table 18. Type of attack encountered depending on the cyber maturity

Type of Attack *q*	Probability
DDoS	29.6%
Ransomware	17.6%
Loss of data (other than ransomware)	9.8%
Business email compromise	30.8%
Other	12.2%

Regarding the probabilities in Table 18, let us note that no effect of correlation between these probabilities and the cyber maturity has been identified until now. Adding such an effect should require more information about the hackers' strategy and how they perceive (and adapt) to the cyber maturity of the potential victim.

Table 19. Probability that each point of entry is involved. In the case of a phishing email, in 51% of cases, it is associated with credential theft. Credential theft with no phishing has a probability of 29%.

Entry point *q*	Probability
Phishing email	63%
Credential theft	44%
Third party	40%
Unpatched server	28%
Brute force server credential	17%

A.2.3. Parameters Related to the Severity of the Attack

Table 20. Probability of defending against the attack

Maturity	Probability of Defending
Novice	7.5%
Intermediate	9.5%
Expert	11.0%

Table 21. Probability of implementing a backup plan as a function of the number of days since the start of the outage, for companies with turnovers of less than $1 billion

Day	1	2	3	4	5	6	7	None/fail
Probability	0.27	0.08	0.05	0.04	0.03	0.02	0.01	0.50

Table 22. Probability of implementing a backup plan as a function of the number of days since the start of the outage, for companies with turnovers of more than $1 billion

Day	1	2	3	4	5	6	7	None/fail
Probability	0.44	0.13	0.08	0.06	0.04	0.03	0.02	0.20

Table 23. Probability distribution governing the time to reach zero (C)BI losses once service to the cloud has been recovered

Day	1	2	3	4	5	6	7
Probability	0.39	0.19	0.13	0.1	0.08	0.06	0.05