Two simple models of business interruption accumulation risk in cyber insurance

Ulrik Franke
RISE Research Institutes of Sweden
Kista, Sweden
ulrik.franke@ri.se

Joachim Draeger
IABG mbH
Ottobrunn, Germany
draeger@iabg.de

***Abstract*—As modern society becomes ever more dependent on IT services, risk management of cyber incidents becomes more important. Cyber insurance is one tool, among others, for such risk management that has received much attention in the past few years. One obstacle to well-functioning cyber insurance, however, is the fact that cyber accumulation risk remains poorly understood, despite efforts from practitioners and scientists. In this article, we address the accumulation risk of business interruption incidents, an area that has received less attention than the accumulation risk of data breach incidents.**

Two simple models are introduced: First, a model that takes the insurer's perspective and explores the impact on aggregated claims cost from incidents that unintentionally propagate between firms. Second, a model that takes the insured's perspective, considering the impacts of limited incident management capacity and showing that there is sometimes an economic case for collectively funding additional incident managers. The paper is concluded with some reflections on the models and an outlook.

***Index Terms*—cyber insurance, business interruption, accumulation risk**

I. Introduction

Cyber incidents, whether antagonistic attacks or non-antagonistic disruptions resulting from mistakes or mismanagement, are never fully unavoidable. Accordingly, a firm may have to recover from an uncertain number of cyber related incidents each year. Since the efforts to deal with these incidents can hardly be predicted in advance, cyber insurance is gaining attention as one cyber risk management tool among others. Not only insurance professionals and academics, but also national governments like the the UK [1] and Singapore [2], and international organizations like the EU Agency for Network and Information Security (ENISA) [3], [4] and the OECD [5] have taken in interest in cyber insurance lately.

While in principle cyber insurance is thus an interesting tool for cyber risk management, there are also a number of obstacles. On the supply side, the OECD identifies problems of (i) risk quantifiability, (ii) acculumation risk, and (iii) reinsurance availability [5, pp. 94–101]. On the demand side, the risks of (iv) lack of awareness of potential losses, (v) misunderstandings about coverage, and (vi) unsuitability of the coverage available [5, pp. 101–104] were identified.

In this article, we focus on accumulation risk, perhaps the single cyber insurance obstacle that has received the most attention. For example, the Geneva Association, self-proclaimed leading international think tank of the insurance industry, recently described accumulation risk as being "at the heart of many concerns about cyber risk".[1]

While the importance of cyber acculumation risk is thus widely acknowledged, not all aspects of it get equal attention. In particular, data breach and related third-party liabilities serve as the point of reference in most discussions. A partial explanation for this focus is that third-party liabilities as defined by breach notification laws, enacted in most US states, have been an important driver for the growth of the US cyber insurance market [6], which accounts for over 85 % of global cyber insurance volume [7]. However, cyber insurances also cover business interruptions and related first-party costs, an aspect that has traditionally been more emphasized on the European cyber insurance market [6], at least before the General Data Protection Regulation (GDPR).

Asking whether data breach or business interruption is the most important is a false start. This cannot be determined in the abstract – it depends on many factors, including the kind of business that is insured, the details of the insurance policy, the technology used, and not least the incident scenario envisioned. However, it is safe to say that cyber accumulation risk applies to business interruption no less than data breach, and that the costs involved can be considerable. A 2018 study by Lloyd's, considering the consequences of service outages at major cloud service providers (Google, Amazon, Microsoft, or IBM) explores business interruption scenarios of some 3-6 days, which would result in ground-up losses between $6.9 and $14.7 billion and between $1.5 and $2.8 billion in industry insured losses in the US alone [8]. An even more recent study from Lloyd's explores the global impacts of contagious malware, similar to the NotPetya or WannaCry cases [9]. In all three variants of this scenario, business interruption represents the greatest insured losses by far, being several times greater than incident response costs or liabilities [9, Table 9, p. 41].

It is against this background that we study the cyber accumulation risk aspects of business interruption. More precisely, we first address accumulation risk from the insurers perspective, with a traditional monetary focus.The novel contribution

U. Franke was partially supported by the Swedish Civil Contingencies Agency, MSB (agreement no. 2015-6986).

[1]The Geneva Association, *Addressing cyber accumulation risk*, August 16, 2018, https://www.genevaassociation.org/addressing-cyber-accumulation-risk, accessed February 25, 2019.

here is a model where Poisson incidents unintentionally *propagate* from one firm to another, allowing incident numbers to be represented by independent superimposed Poisson processes, even though the resulting aggregated claims cost exhibit the larger variance typical of dependent incidents.

Secondly, we take the insureds perspective, considering the impacts of limited incident management capacity. Most cyber insurance policies include not only monetary compensation, but also incident management support, typically with a one-stop-shop incident telephone service [6]. The first response is often coordinated by a law firm that calls upon IT consultancies, other law firms, PR consultants etc. as needed. As pointed out by a reinsurer in the interviews done by Franke [6], non-monetary resources such as consultants could then become bottlenecks in incident handling if many incidents, unexpectedly, occur simultaneously. The novel contribution here is a simple model that captures this phenomenon.

The remainder of this paper is structured as follows. After some related work in Section II, we first introduce some preliminaries and a baseline model in Section III. Subsequently, the consequences of incident propagation between firms (Section IV) and limited incident capacity management (Section V) are analyzed. Section VI concludes the paper.

II. Related Work

There is much literature on cyber insurance. A good but now slightly dated literature review was conducted by Böhme & Schwartz [10]. More recent reviews include Eling & Schnell [11] and Marotta et al. [12].

The importance of treating cyber accumulation risk is emphasized in virtually all the literature on cyber insurance, and indeed information security economics more generally. Thus, Anderson & Moore in their seminal *Science* article from 2006 point to two specific cyber risk interdependencies that are problematic from an insurance perspective [13]: (i) that the IT of one firm is connected to others, meaning that incidents can spread, and (ii) that many firms use the same IT product, meaning that the same vulnerabilities are found in many places. Hence, in recent years much effort has been directed to the investigation of cyber risk correlations, interdependencies and accumulation. Nevertheless, the field remains full of open questions. Eling & Schnell, summarizing their review, still identify the issues of modeling and aggregation of cyber risk as areas in need of more future work [11].

However, most of the cyber risks literature focuses on *data breaches*, aiming, e.g., to find models for the frequency (beta-binomial [14], Poisson or negative binomial [15], negative binomial [16]) and severity (power law [17], log-normal [16], log-skew-normal distribution [15] truncated Pareto [18]) of breaches. By contrast, the properties of *business interruption* incidents, including accumulation risks, have not received as much attention in the cyber insurance literature. This is where our paper aims to make a contribution, proposing two simple models for reasoning about accumulation risk related to business interruption.

III. Preliminaries

In the following, we define the basic building blocks of the models that are developed in Sections IV and V. Let a firm j experience a random number N of IT service outages, where N has a Poisson distribution with mean λ:

$$\mathsf{Pr}(N(t)=k)=\frac{e^{-\lambda t}(\lambda t)^k}{k!} \tag{1}$$

We will typically assume that the time period $[0,t]$ is a year, so that $t=1$ and can be omitted.

Once an IT service outage occurs, let its duration T be log-normally distributed with parameters a and σ, i.e., with the following probability density function:

$$f_T(x)=\frac{1}{x\sigma\sqrt{2\pi}}\exp\left(-\frac{(\ln(x)-a)^2}{2\sigma^2}\right) \tag{2}$$

Empirical studies have shown the lognormal distribution to be suitable for the durations of IT service outages [19], [20].

For simplicity, let the downtime cost $C(T)$ be linear, $C(T)=c\cdot T$. Thus, in a given year, the firm's *aggregate downtime cost* S_j is a random sum of N independent random variables. The probability distribution of S is the convolution of the two underlying distributions; the Poisson distribution for the count of incidents and the lognormal for the size of them. In general, such convolutions are complicated, but for the special case that N is Poisson distributed, the aggregate distribution is a *compound Poisson distribution* where the mean and variance take the following simple forms [21, p. 110]:

$$\mathsf{E}[S_j]=\lambda\cdot\mathsf{E}[C(T)]=\lambda\cdot c\cdot\mu^{(1)} \tag{3}$$

$$\mathsf{Var}[S_j]=\lambda\cdot\mathsf{E}[(C(T))^2]=\lambda\cdot c^2\cdot\mu^{(2)} \tag{4}$$

where $\mu^{(k)}=\mathsf{E}[T^k]$ are the raw moments of the distribution of T. For the lognormal distribution, we have:

$$\mu^{(k)}=\exp\left(ka+\frac{k^2\sigma^2}{2}\right) \tag{5}$$

Now, let an insurer insure m firms of the kind defined above (insureds). For simplicity, we omit deductibles and limits. The insured j suffers from $N_j(\lambda)$ annual insured cyber incidents, with durations that are realizations of T_j. Since the insureds are identical but independent the subscripts j can be omitted and the mean and variance of the entire portfolio of m insureds are just the sums of those of the individual firms, as independence makes all covariances zero:

$$\mathsf{E}[S]=\sum_{j=1}^{m}\mathsf{E}[N_j]\cdot\mathsf{E}[C(T_j)]=m\cdot\lambda\cdot c\cdot\mu^{(1)} \tag{6}$$

$$\mathsf{Var}[S]=\sum_{j=1}^{m}\mathsf{E}[N_j]\cdot\mathsf{E}[(C(T))^2]=m\cdot\lambda\cdot c^2\cdot\mu^{(2)} \tag{7}$$

As an insurer accepts the obligation to indemnify the stochastic loss sum S, it also accepts the *risk* that this loss deviates from its expectation. There are several, related, measures of this risk in the literature [22, pp. 216–219].

The basic question of risk in the context of a portfolio of insureds is: What is the risk that the arithmetic mean $\bar{s} = \sum_{i=1}^{m} \frac{s_i}{m}$ of multiple realizations of S exceed the expectation $\mathsf{E}[S]/m$? From the law of large numbers, we know that for a loss sum S with finite expectation and variance, this probability tends to zero in the limit of an infinite pool of insureds:

$$\lim_{m\to\infty} \Pr\left(\left|\bar{s} - \frac{\mathsf{E}[S]}{m}\right| > \varepsilon\right) = 0 \tag{8}$$

When *not* in the limit, of course, this behavior depends on the variance of $\bar{s}$:

$$\mathsf{Var}[\bar{s}] = \mathsf{Var}\left[\frac{S}{m}\right] = \frac{\mathsf{Var}[S]}{m^2} = \tag{9A}$$

$$= \frac{m\sigma^2}{m^2} = \frac{\sigma^2}{m} \tag{9B}$$

Here, (9B) applies when the claims are independent, all with the same finite variance σ^2. If claims are not independent, or of they are independent, but with different variances, (9A) must be used instead.

Letting $\varepsilon = k\sqrt{\mathsf{Var}[\bar{s}]}$ in (8), we obtain:

$$\lim_{m\to\infty} \Pr\left(\left|\bar{s} - \frac{\mathsf{E}[S]}{m}\right| > k\sqrt{\mathsf{Var}[\bar{s}]}\right) =$$

$$= \lim_{m\to\infty} \Pr\left(\left|\bar{s} - \frac{\mathsf{E}[S]}{m}\right| > k\frac{\sqrt{\mathsf{Var}[S]}}{m}\right) = 0 \tag{10A}$$

$$= \lim_{m\to\infty} \Pr\left(|\bar{s} - \mu| > k\frac{\sigma}{\sqrt{m}}\right) = 0 \tag{10B}$$

As before, (10B) applies when claims are independent with the same finite expectation μ and the same finite variance σ^2. Otherwise, (10A) must be used. This leads to the first version of insurer's relative risk [22, p. 217]:

Definition 1. Insurer's relative risk (IRR_1) consists in the possibility that the mean loss exceeds its expected value by more than k-fold of its standard deviation.

Dividing (12) through by $\mathsf{E}[S]/m$ gives a normalized version:

$$\lim_{m\to\infty} \Pr\left(\left|\frac{\bar{s} - \frac{\mathsf{E}[S]}{m}}{\frac{\mathsf{E}[S]}{m}}\right| > k\frac{\sqrt{\mathsf{Var}[S]}}{\mathsf{E}[S]}\right) = 0 \tag{11A}$$

$$= \lim_{m\to\infty} \Pr\left(\left|\frac{\bar{s} - \mu}{\mu}\right| > k\frac{\sigma}{\mu\sqrt{m}}\right) = 0 \tag{11B}$$

As before, (11B) applies when claims are independent with the same finite expectation μ and the same finite variance σ^2. Otherwise, (11A) must be used. This leads to the second version of *insurer's relative risk* [22, p. 217]:

Definition 2. Insurer's relative risk (IRR_2) consists in the possibility that the relative deviation of mean loss exceeds the k-fold of the coefficient of variation ($\sqrt{m}\mathsf{Var}[S]/\mathsf{E}[S]$ or, when applicable, simply σ/μ).

Following [22], we set $k = 1$, and use the right-hand-side of the inequality in (11A), i.e., $\sqrt{\mathsf{Var}[S]}/\mathsf{E}[S]$, as the numerical IRR_2 risk measure. The inverse dependence on the square root of m entails that increasing the pool of insured decreases the risk, as illustrated in the following example.

Example 1. Let an insurer insure m identical but independent customers, each of which has an expected annual outage cost of $\lambda \cdot c \cdot \mu^{(1)} = \$22\mathrm{k}$ with a standard deviation of $\sqrt{\lambda \cdot c^2 \cdot \mu^{(2)}} = \$15\mathrm{k}$. The insurer's relative risk IRR_2 for a selection of different m values is given in Table I.

TABLE I
INSURER'S RELATIVE RISK IRR_2 BY NUMBER OF INSUREDS m.

m	IRR_2
1	0.6818
10	0.2156
100	0.0682
1 000	0.0216
10 000	0.0068
100 000	0.0022

In this case, IRR_2 can be calculated by either (11A) or (11B), as claims are independent. Note how increasing m by a factor of 100 decreases the risk by a factor of $\sqrt{100} = 10$. This strategy to decrease risk is a feature of independent claims. In the next section, this feature will change.

IV. Incident propagation

In the preceding section, we defined a model where each firm has its private frequency of incidents, following a Poisson distribution, and its private incident durations, following a log-normal distribution. Let us first examine the Poisson distribution of incidents in firm j, characterized by intensity λ_j. The incidents lumped together in λ_j will typically emanate from different systems. For example, in one firm, some incidents will come from the CRM system, some from the ERP system, some from the e-mail system, some from the SCADA system, etc. Since the union of several Poisson processes is itself a Poisson process with an intensity corresponding to the sum of the intensities of the constituent processes, it is reasonable to perceive λ_j as such a sum:

$$\lambda_j = \lambda_{\mathrm{CRM}_j} + \lambda_{\mathrm{ERP}_j} + \lambda_{\mathrm{e-mail}_j} + \lambda_{\mathrm{SCADA}_j} + \dots$$

This model tacitly corresponds to a silo model of a firm running its own IT, unconnected to others. However, such independence is rarely the case. Nowadays, the security of any single firm is highly dependent on the security of other firms with whom there is electronic interaction, i.e., virtually any exchange of goods or services. Incidents can then propagate through electronic supply chains between firms: If firm B depends on firm A and firm A is down due to an incident, firm B will be down as well. As explained in the introduction, the impact of one-to-many dependencies, such as many firms using a single cloud service provider, are especially worrisome. In the following, we use the adoption of cloud services for expositional purposes, but the model developed is applicable to a wider range of situations.

As an example but without loss of generality, consider m firms adopting a cloud CRM solution. Their respective Poisson

TABLE II
INSURERS RELATIVE RISK IRR_2 BY NUMBER OF INSUREDS m AND PROPAGATION COEFFICIENT α.

m	$\alpha = 0$	0.01	0.05	0.1	0.2	1
1	0.6818	0.6818	0.6818	0.6818	0.6818	0.6818
10	0.2156	0.2251	0.2596	0.2972	0.3608	0.6818
100	0.0682	0.0962	0.1663	0.2251	0.3109	0.6818
1 000	0.0216	0.0715	0.1539	0.2166	0.3055	0.6818
10 000	0.0068	0.0685	0.1526	0.2157	0.3050	0.6818
100 000	0.0022	0.0682	0.1525	0.2156	0.3049	0.6818

processes now get common terms for the CRM incident process:

$$\begin{aligned}\lambda_1 &= \lambda_{\text{CRM}} + \lambda_{\text{ERP}_1} + \lambda_{\text{e-mail}_1} + \lambda_{\text{SCADA}_1} + \dots \\ \lambda_2 &= \lambda_{\text{CRM}} + \lambda_{\text{ERP}_2} + \lambda_{\text{e-mail}_2} + \lambda_{\text{SCADA}_2} + \dots \\ &\vdots \\ \lambda_m &= \lambda_{\text{CRM}} + \lambda_{\text{ERP}_m} + \lambda_{\text{e-mail}_m} + \lambda_{\text{SCADA}_m} + \dots\end{aligned}$$

Such common terms have consequences for the aggregated characteristics of the incidents.

Let m firms with independent and identically distributed incident processes (all $\lambda_j = \lambda$) all adopt the same cloud-based CRM system. Assume for simplicity and clarity of the example that the new system has the same expected number of incidents as the old ones and that these CRM incidents account for a fraction $\alpha \in [0, 1]$ of all the incidents of each firm, i.e., $\lambda_{\text{CRM}} = \alpha\lambda$. Ex ante, the superpositioned Poisson process was the sum of m independent processes:

$$\mathsf{E}[N] = \sum_{j=1}^{m} \lambda_j = m\lambda \qquad \mathsf{Var}[N] = \sum_{j=1}^{m} \lambda_j = m\lambda \tag{12}$$

Ex post, we still superimpose independent Poisson processes, so the superposition is also Poisson, but its components are different:

$$\begin{aligned}\mathsf{E}[N] &= \sum_{j=1}^{m}(1-\alpha)\lambda_j + \lambda_{\text{CRM}} = \\ &= (m(1-\alpha)+\alpha)\lambda \le m\lambda \qquad (13) \\ \mathsf{Var}[N] &= \sum_{j=1}^{m}(1-\alpha)\lambda_j + \lambda_{\text{CRM}} = \\ &= (m(1-\alpha)+\alpha)\lambda \le m\lambda \qquad (14)\end{aligned}$$

Thus, counting the aggregate *number* of incidents, both the expectation and the variance decrease.

However, whereas an incident in the previous, individual, CRM systems affected just a single firm, an incident in the new, cloud-based, CRM system affects all m firms. The difference becomes clear when modifying (6) and (7) accordingly:

$$\begin{aligned}\mathsf{E}[S] &= \sum_{j=1}^{m}(1-\alpha)\lambda_j \cdot \mathsf{E}[C(T)] + \\ &\quad + \lambda_{\text{CRM}} \cdot \mathsf{E}[mC(T)] = \\ &= (m(1-\alpha)+m\alpha)\lambda \cdot c \cdot \mu^{(1)} = \\ &= m \cdot \lambda \cdot c \cdot \mu^{(1)} \qquad (15) \\ \mathsf{Var}[S] &= \sum_{j=1}^{m}(1-\alpha)\lambda_j \cdot \mathsf{E}[(C(T))^2] + \\ &\quad + \lambda_{\text{CRM}} \cdot \mathsf{E}[(mC(T))^2] = \\ &= (m(1-\alpha)+\alpha m^2)\lambda \cdot c^2 \cdot \mu^{(2)} \qquad (16)\end{aligned}$$

Here, a fraction $1 - \alpha$ of all incidents occur isolated at a single firm and incur the individual cost $C(T_n) = c \cdot T_n$. The remaining fraction α of incidents occur in the cloud, affects all firms, and incurs the collective cost $mC(T_n) = mc \cdot T_n$. For the expectation, this does not matter – α is cancelled out in (15). But for the variance, it matters more, as the linear dependence on m is replaced by a quadratic one in (16) as α grows.

Example 2. Continuing Example 1, we analyze the effect of α on the insurer's relative risk with m customers, each of which has an expected annual outage cost of \$22k with a standard deviation of \$15k. The first column in Table II, where $\alpha = 0$, is identical to Table I. However, with an increasing α, the risk grows.

For $\alpha = 1$, there is no reduction of relative risk, regardless of the number of insureds. In this degenerate case, all outage occurrences come from a single Poisson process, just as for a single insured. Thus, all values in the first row ($m = 1$), and the last column ($\alpha = 1$) are identical.

It is interesting to note how even a small α offsets the risk decrease of a large m. For example, with $\alpha = 0.01$ it takes 100 000 insureds to reach the same risk as for 100 insureds with $\alpha = 0$. With $\alpha = 0.1$, 100 000 insureds entail as much risk as just 10 insureds with $\alpha = 0$. The effects of an increasing α are thus almost impossible to offset by increasing m.

Remark. The model expressed in (15) and (16) is of course simplified in many ways, especially in that all firms are homogeneous, sharing the same λ, c, α, $\mu^{(1)}$, and $\mu^{(2)}$ parameters. However, these parameters can be seen as *average*

values of individual parameters, e.g., α being the average of m different fractions of incidents caused by systems that are replaced by solutions that allow for incident propagation. Individual parameters would make the sums in Eqs. (15) and (16) much longer, but not change the main observation about growing variance and risk. Indeed, underestimating the parameter variability just makes the model conservative.

Remark. For business interruption events, Franke found empirical evidence that insurers do apply strategies to control α [6]. One of the insurers interviewed offers coverage of outages caused by external service providers either (i) for any provider with an increase in the premium of some 20-25% and the indemnity limit cut in half, or (ii) for some 3-5 specific providers, named ex ante.

V. Limited incident management capacity

As before, let a firm have a cyber insurance that reimburses the cost of breakdown without deductibles or limits. The actuarially fair insurance premium is just the expected value $\mathsf{E}[S] = \lambda \cdot c \cdot \mu^{(1)}$.

However, in some situations, incident handling may not be instantaneous. Suppose that an external incident management service is needed to restore service. When this service is immediately available, the duration of the breakdown is as above, with expectation $\mu^{(1)}$. If the service is *not* immediately available, however, the breakdown is prolonged due to the queueing situation.

It is not difficult to understand why this could quite often be the case. At market equilibrium, the supply of incident managers corresponds to the demand, perceived as an average over a relatively short time interval. But occurrences of incidents where many firms simultaneously experience outages may be so rare that they are absent from the sampling window. If the number of incident managers corresponded to the demand in such extreme events, many of them would sit idle in normal circumstances. So as soon as incidents are clustered for whatever reason, the few available incident managers become a scarce resource.

Specifically, let m identical but independent firms be served by a single incident management service provider, which is sufficient at market equilibrium and at steady-state of the resulting queue. Now, if an incident occurs so that all m firms are hit at the same time, the expected outage duration is no longer just $\mu^{(1)}$. In the absence of a priority order, there is an equal probability for any one firm ending up first, second, …, mth in the queue. The total time from incident beginning to service restoration for any one firm is then the sum of a random number, uniformly distributed over $\{0, \ldots, m-1\}$, of others' service restorations plus a single own service restoration. The expected value of this convolution is simply the product of the respective expected values of the two distributions: [21, p. 109]:

$$\mathsf{E}[T_{\text{total}}] = \left(\frac{m-1}{2} + 1\right)\mu^{(1)} \tag{17}$$

Remark. This model is applicable if no new incidents arrive during service – that would require a full birth-death queue model. This is realistic if incidents are rare and service times comparably short. While minor service outages happen all the time, major service outages are indeed quite rare – they might happen a few times annually and are typically solved in the order of hours or, worst case, days.

Remark. (17) might be too pessimistic. First, it could be that separate incident management for each firm is unnecessary – once service is restored centrally, e.g., in a cloud service, service is fully restored everywhere. However, even if this is true on a technical level for the original service, there might still be additional residual cascade effects to manage, e.g., other technical systems affected, backlogs of waiting customers or activities, legal or PR activities, etc. Another reason for more optimism is *learning* – even if incident management for each firm is necessary, it might be that the effects are easier and quicker to manage the mth time compared to the first. It is straightforward to see that even a modest such learning could have a substantial impact. Let r be a discount factor such that if the expected incident management duration of the first incident managed is $\mu^{(1)}$, that of the second incident managed is $r\mu^{(1)}$ and that of the mth incident managed is $r^{m-1}\mu^{(1)}$. The expected time to serve all m firms is then a geometric sum of m progressively shorter expected service times:

$$\mathsf{E}[T_{\text{all firms}}] = \sum_{i=0}^{m-1} r^i \mu^{(1)} = \frac{1-r^m}{1-r}\mu^{(1)} \tag{18}$$

Even a small learning can have a large effect – for $m = 100$, $r = 0.95$ reduces the expected service time of all m firms from $100\mu^{(1)}$ to approximately $20\mu^{(1)}$. In such situations, queues will be less of a problem. On the other hand, there are also cases where learning might be less effective, e.g., a cryptovirus hitting several interdependent firms at the same time. Here, the resolution might require individual restoration from backups, roll-backs, legal and PR activities, etc. where there is little scope for learning. Thus, while many different cases are discernable, the model introduced here is intended to reflect cases where (i) individual incident management is needed and (ii) there is little scope for learning.

Under the naive assumption of cyber insurance without limits on claims, the insured firm would get compensated for the entire duration of an outage. However, in practice, there are always limits to the policies, and such limits might well be reached while waiting for the incident management service.

This suggests an alternative risk-management strategy, a complement to insurance, viz. to collectively fund more incident management service providers. Assume for simplicity that this can be achieved at no coordination cost. Each firm then faces the following problem of minimizing its expected cost:

$$\min_p \lambda \cdot c \cdot \mu^{(1)} \left(\frac{m-1}{2p} + 1\right) + \frac{p-1}{m} c_p \tag{19}$$

Here, p is the number of service providers and c_p is the cost needed to subsidize each additional service providers beyond the first one. The first term in the total cost is the downtime cost as identified above, but the waiting time in the queue is

split among the p service providers. The second term is the subsidy of $p-1$ service providers split over the m firms.

The optimal number of service providers p^* is found by taking the derivative with respect to p and requiring it to be zero:

$$\frac{c_p}{m} - \lambda \cdot c \cdot \mu^{(1)} \left(\frac{m-1}{2p^{*2}} \right) = 0 \tag{20}$$

Rearranging yields:

$$p^* = \sqrt{m(m-1)\frac{\lambda \cdot c \cdot \mu^{(1)}}{2c_p}} \tag{21}$$

Unsurprisingly, the optimal number of service providers increases with the number m of firms affected and with the annual expected cost $\lambda \cdot c \cdot \mu^{(1)}$ (without incident management queuing) of outages. It decreases with the cost c_p of operating service providers. For large m, p^* approaches linear growth with respect to m.

Whenever $p^* > 1$, it pays to first subsidize more incident management service providers and only then insure the residual outage cost at an actuarially fair premium, compared to just insuring the outage cost in the absence of more incident management service providers at an actuarially fair premium.

Example 3. Say that the annual subsidy cost c_p for an incident management service provider is \$100 k. Then a mere three firms with expected annual outage costs of \$134 k in the absence of queues would find it worthwhile to subsidize an additional provider ($p^* = 2$). Even if the expected expected annual outage costs in the absence of queues are just \$9 k, ten firms would still find it worthwhile to subsidize an additional provider ($p^* = 2$).

However, for calculations such as those in Example 3 to hold, the queue resulting from limited incident capacity management must occur for *every incident*. But if such incidents are frequent enough, this would shift the market equilibrium towards a greater number of incident managers. Thus, Example 3 is only realistic for very small λ values, i.e., even though there is a queue for *every* incident, incidents are still so *rare* that service providers will sit idle and supply and demand do not match.

A more common case is that just a fraction $\beta \in [0,1]$ of business interruption incidents lead to the queuing situation. Each firm then faces a modified minimization problem:

$$\begin{aligned} \min_p &(1-\beta)\lambda \cdot c \cdot \mu^{(1)} + \\ &+ \beta \cdot \lambda \cdot c \cdot \mu^{(1)} \left(\frac{m-1}{2p} + 1 \right) + \frac{p-1}{m} c_p \end{aligned} \tag{22}$$

Since the cost for an outage in the absence of queues (the first term in (22)) is not dependent on p and thus disappears in the differentiation, the optimal number of service providers is just (21) modified by $\sqrt{\beta}$:

$$p^* = \sqrt{m(m-1)\frac{\beta \cdot \lambda \cdot c \cdot \mu^{(1)}}{2c_p}} \tag{23}$$

Remark. Following Section IV, it is reasonable to wonder about the relationship between β in (22)–(23) and α in (15)–(16). Could they be identical? Indeed they could: if incidents propagate between m firms according to α, and these same m firms then require simultaneous individual incident management services for these incidents but not others, then $\alpha = \beta$. However, equality could fail: If the propagation set and queuing set of firms are not identical, if not all incidents propagated require individual incident management, or if not all simultaneous incidents that require individual incident management originate from propagation, then α and β are not necessarily identical.

For small β, the number of firms that need to affected for the subsidy to be worthwhile grows to the point where the assumption of no coordination cost is no longer credible. However, a diligent insurer might take the role of broker. Say that the insurer is aware that it has many customers in its portfolio that are all dependent on the same kind of technology and would be stuck in a queue waiting for an incident management service provider if a common cause error struck them all at the same time. Instead of charging a correspondingly higher (actuarially fair) premium or a lower (actuarially fair) premium with a limit that would leave the insureds without compensation for most of a prolonged outage, the insurer could add the per-customer subsidy $\frac{p-1}{m} c_p$ to the premium of each customer and make the subsidized incident management service providers available to its m customers.

Example 4. Let there be $m = 1\,000$ interdependent customers that run a risk $\beta = 1\%$ of simultaneous need for incident management. Let each firm have an expected annual outage cost of \$20 k in the absence of queues. As in Example 3, $c_p =$ \$100 k. With just a single service provider, however, the 1% chance of the queue would increase the actuarially fair premium (or, equivalently, the firm's expected cost in the absence of insurance) to 99% times the \$20 k plus 1% times 500.5 times the \$20 k, i.e., a premium of \$120 k. However, $p^* \approx 32$, so the insurer could subsidize 31 additional service providers and insure the residual risk, all at a premium of approximately \$26 k per insured firm.

VI. Conclusions and Outlook

In this paper, we have addressed an underinvestigated aspect of the well-known problem of cyber accumulation risk: the accumulation risk of business interruption incidents. The models introduced in Sections IV and V are simple, but still capture some of the relevant dynamics of these phenomena. Section IV adopts the insurer's perspective and a traditional monetary focus, whereas Section V adopts the insured's perspective, showing that incident management capacity might become a scarce resource.

While our models could be improved in many ways, leading to more accurate representations of empirical circumstances and also to more profound theoretical insights, we believe that the simplicity also has some virtues, shedding light on some basic phenomena related to the accumulation risk

of business interruption incidents. The incident propagation model illustrates the limitations to insurer risk management by simple portfolio growth: the numbers needed to offset even modest incident propagation are prohibitive. While this insight is not new, it is useful to have a simple and illustrative model that can be used as a pedagogical tool. The limited incident management capacity model is more novel in the sense that it represents, as far as we know, the first model of this phenomenon.

To conclude, the accumulation risk of business interruption incidents continues to be an interesting area for future research. Areas worthy of investigation include but are not limited to (i) the relationship between business interruption and data breach accumulation risk, (ii) incentives and interplay between insurers and insureds, (iii) empirical studies of business interruption incidents that unintentionally propagated between firms, (iv) empirical studies of limited incident management capacity, and (v) reinsurance aspects of business interruption accumulation risk.

ACKNOWLEDGMENT

The authors are grateful for comments on the manuscript by Dr. Pontus Svenson of Research Institutes of Sweden. U. Franke also wishes to acknowledge discussions on cyber incident propagation with Prof. Shaun Wang of Nanyang Technological University, Singapore.

REFERENCES

[1] Cabinet Office, "Cyber Insurance Market: Joint Government and Industry Statement," November 5 2014, accessed on January 9, 2017. [Online]. Available: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/371036/Cyber_Insurance_Joint_Statement_5_November_2014.pdf

[2] Monetary Authority of Singapore, "A Bold Approach to Cyber Risk Management – Opening Address by Mr Bernard Wee, Executive Director, Monetary Authority of Singapore, at the Asia Cyber Risk Summit," May 16 2016, accessed on May 30, 2018. [Online]. Available: http://www.mas.gov.sg/News-and-Publications/Speeches-and-Monetary-Policy-Statements/Speeches/2016/A-Bold-Approach-to-Cyber-Risk-Management.aspx

[3] ENISA, "Cyber insurance: Recent advances, good practices and challenges," European Union Agency for Network and Information Security, Tech. Rep., 2016. [Online]. Available: http://dx.doi.org/10.2824/065381

[4] ——, "Commonality of risk assessment language in cyber insurance," European Union Agency for Network and Information Security, Tech. Rep., 2017.

[5] OECD, "Enhancing the Role of Insurance in Cyber Risk Management," 2017.

[6] U. Franke, "The cyber insurance market in Sweden," *Computers & Security*, vol. 68, pp. 130–144, 2017.

[7] D. M. Hofmann, S. Wilson, and R. A. Carter, "Advancing accumulation risk management in cyber insurance," The Geneva Association, Tech. Rep., 2018, accessed February 25, 2019. [Online]. Available: https://www.genevaassociation.org/research-topics/cyber-and-innovation/advancing-accumulation-risk-management-cyber-insurance

[8] "Cloud Down: Impacts on the US economy," Lloyd's of London, Tech. Rep., 2018, accessed March 19, 2018. [Online]. Available: https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/cloud-down

[9] J. Daffron, S. Ruffle, C. Andrew, J. Copic, K. Quantrill, S. A., and E. Leverett, "Bashe attack: Global infection by contagious malware," Cambridge Centre for Risk Studies, Lloyds of London and Nanyang Technological University, Tech. Rep., 2019, accessed February 4, 2019. [Online]. Available: https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/bashe-attack

[10] R. Böhme and G. Schwartz, "Modeling Cyber-Insurance: Towards a Unifying Framework." in *Workshop on Economics of Information Security – WEIS*, 2010.

[11] M. Eling and W. Schnell, "What do we know about cyber risk and cyber risk insurance?" *The Journal of Risk Finance*, vol. 17, no. 5, pp. 474–491, 2016. [Online]. Available: http://dx.doi.org/10.1108/JRF-09-2016-0122

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

[13] R. Anderson and T. Moore, "The economics of information security," *Science*, vol. 314, no. 5799, pp. 610–613, 2006. [Online]. Available: http://dx.doi.org/10.1126/science.1130992

[14] R. Böhme and G. Kataria, "Models and measures for correlation in cyber-insurance," in *Workshop on Economics of Information Security – WEIS*, 2006.

[15] M. Eling and N. Loperfido, "Data breaches: Goodness of fit, pricing, and risk measurement," *Insurance: mathematics and economics*, vol. 75, pp. 126–136, 2017.

[16] B. Edwards, S. Hofmeyr, and S. Forrest, "Hype and heavy tails: A closer look at data breaches," *Journal of Cybersecurity*, vol. 2, no. 1, pp. 3–14, 2016.

[17] T. Maillart and D. Sornette, "Heavy-tailed distribution of cyber-risks," *The European Physical Journal B*, vol. 75, no. 3, pp. 357–364, 2010.

[18] S. Wheatley, T. Maillart, and D. Sornette, "The extreme risk of personal data breaches and the erosion of privacy," *The European Physical Journal B*, vol. 89, no. 1, p. 7, 2016.

[19] U. Franke, H. Holm, and J. König, "The distribution of time to recovery of enterprise IT services," *IEEE Transactions on Reliability*, vol. 63, no. 4, pp. 858–867, December 2014.

[20] B. Schroeder and G. Gibson, "A large-scale study of failures in high-performance computing systems," *IEEE Transactions on Dependable and Secure Computing*, vol. 7, no. 4, pp. 337–350, 2010.

[21] D. Bahnemann, *Distributions for Actuaries*, ser. CAS Monograph Series. Casualty Actuarial Society, 2015, no. 2.

[22] P. Zweifel and R. Eisen, *Insurance economics*. Springer Science & Business Media, 2012.