Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models

Caroline Hillairet, Olivier Lopez

► To cite this version:

Caroline Hillairet, Olivier Lopez. Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models. 2020. hal-02564462v1

HAL Id: hal-02564462

<https://hal.science/hal-02564462v1>

Preprint submitted on 5 May 2020 (v1), last revised 30 Nov 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models.

Caroline HILLAIRET¹, Olivier LOPEZ².

May 5, 2020

Abstract

In this paper, we propose a general framework to design accumulation scenarios that can be used to anticipate the impact of a massive cyber attack on an insurance portfolio. The aim is also to emphasize the role of countermeasures in stopping the spread of the attack over the portfolio, and to quantify the benefits of implementing such strategies of response. Our approach consists of separating the global dynamic of the cyber event (that can be described through compartmental epidemiological models), the effect on the portfolio, and the response strategy. This general framework allows us to obtain Gaussian approximations for the corresponding processes, and sharp confidence bounds for the losses. A detailed simulation study, which mimics the effects of a Wannacry scenario, illustrates the practical implementation of the method.

Key words: Cyber insurance; emerging risks; counting processes; compartmental epidemiological models; risk theory.

Short title: Propagation of cyber incidents in an insurance portfolio

¹ Ensaie Paris, Centre de Recherche en Economie et Statistique, CREST, 5 avenue Henry Le Chatelier, F-91120 Palaiseau, France, E-mail: caroline.hillaiRET@ensae.fr

² Sorbonne Université, CNRS, Laboratoire de Probabilités, Statistique et Modélisation, LPSM, 4 place Jussieu, F-75005 Paris, France, E-mail: olivier.lopez@sorbonne-universite.fr

1 Introduction

The growing cyber threat encourages companies into subscribing insurance policies to complete their arsenal of protection. The cyber-insurance market is increasingly developing (see for example Camillo (2017)), with important uncertainties on the real value of the guarantees. Recently, efforts have been made to quantify the risk with a traditional frequency/severity approach, as in Eling and Loperfido (2017), Forrest et al. (2016), Farkas et al. (2019) or Bessy-Roland et al. (2020). Apart from the difficulty to evaluate a fast evolving risk with a relative lack of consistent historical data, the systemic potential of a "cyber hurricane" is a major concern. Indeed, a global failure of the portfolio would break the mutualisation principle at the core of the insurance business. Such type of crisis is referred to as an "accumulation scenario", see Eling and Schnell (2016). The aim of the present paper is to discuss a general framework to design such accumulation scenarios, in order to simulate their impact on an insurance portfolio, depending on the intensity of the attack, and the response of the company and of the policyholders. This methodology may help insurance companies to conceive their strategies in dealing with such cyber hurricanes.

Moreover, we develop an approach which is adapted to the particular nature of the services included in typical cyber contracts, generally a combination of financial repair, and of the fast intervention of expert teams to assist the policyholder in restarting his/her activity. Stamping out the crisis via such assistance is a component of the total cost, and if the company is unable to provide such help, this may increase the total bill considerably. Indeed, if a large number of policyholders are simultaneously victims of an attack, a saturation of this response capacity may occur. Depending on the situation, this may cause an increase of the cost of intervention per policyholder, or an impossibility to honor the contract (by lack of resources that could be mobilized). We propose a way to quantify the probability of such a saturation from a given capacity of response of the insurer. Typically, this requires not only to model the total number of claims during the cyber episode, but also the dynamics of the crisis, by tracking, at each instant, the number of infected policyholder requiring assistance.

The core of our approach is to mix epidemiological models to describe the strength of the attack at a global level, and Gaussian approximation theory to simulate the resulting impact on the insurance portfolio.

Epidemiological models are widely studied in mathematical modeling for biology, to model

the spread of an infectious disease, starting from the standard deterministic SIR model introduced in the seminal papers of McKendrick (1925) and Kermack and McKendrick (1927). While deterministic versions of the SIR model are based on differential equations, stochastic SIR models use Markov chains, branching and diffusion processes. Various extensions have been proposed in the literature, that include vaccination, jumps perturbations, or demographic models. For example, Artalejo and Lopez-Herrero (2014) compute for a stochastic SIR model the whole probability distribution of the time to reach a specific state, and the time to reach a critical number of infections. El Koufi et al. (2019) discuss a stochastic SIR epidemic model with vaccination and environmental fluctuations on the transmission rate, and provide sufficient conditions for the extinction (resp. for the persistence) of the disease. Zhang and Wang (2013) study the asymptotic behavior of a stochastic SIR model with jumps perturbation. Montagnon (2019) points out the role of the population renewal in the persistence of an endemic disease, and thus proposes a model coupling epidemiological multi-type stochastic processes with demographic models. Epidemiological compartmental models have also been used to describe the impact of an epidemic on an insurance portfolio, either using deterministic SIR as in Chen and Cox (2009) or Runhuan Feng (2011), or using a stochastic version as in Lefèvre et al. (2017). Our approach is designed to rely on such compartmental models, but can also be adapted to any other patterns of infection. Indeed, we make the assumption that the contagion does not come from the inside, but more likely from outside the portfolio, since it only gathers a small part of the global affected population. Hence we separately model the global dynamics of the cyber epidemic, and the resulting impact on the portfolio. The framework we develop can be generalized easily to any other type of models for describing the timeline and the strength of the attack.

The rest of the paper is organized as follows. In Section 2, we introduce the general setting that we use to model a cyber event and its result on a portfolio. In Section 3, we use Gaussian approximation theory to approximate the evolution of such episodes for large portfolios, and to derive deviation bounds that help to quantify the probability of saturation of the response. In Section 4, we calibrate a Wannacry-type scenario and discuss the impact of the reaction of the insurance company and of the policyholders through simulation studies.

2 Portfolio

In this section, we introduce a counting process approach to describe the evolution through time of the number of infected policyholders. In Section 2.1, the different variables used to describe the time of infection and the response are described. Section 2.2 is devoted to the description of a SIR (Susceptible - Infected - Recovered) model widely used in epidemiological modeling. Section 2.3 discusses the different measures we use to quantify the impact of a cyber event on the portfolio.

2.1 Modeling the effects of a cyber attack on policyholders

We consider an insurance portfolio composed of n policyholders who are all susceptible of being victim of the contagious cyber-attack. Our approach consists in considering that three types of events may affect a given policyholder :

- infection by the attack;
- if infected, time at which assistance is no longer required;
- immunization by implementing the appropriate patches to the protection strategy.

Each policyholder j will be infected at a different random time T_j (potentially infinite), and then starts the recovery process. Time 0 denotes the origin of the cyber crisis. The attack is stopped for policyholder j at time $T_j + U_j$, where $U_j \geq 0$. The terminology "stopped" may have various meanings depending on the situation and on the structure of the policy. This can simply mean stopping the spread of a malware in the network of the victim, but many cyber policies also include additional help to the policyholder (crisis communication, help to restart the activity of a company...) that can change the duration of the crisis. These examples will be detailed in the application in Section 4.2. The random vector $(T_j, U_j)_{1 \leq j \leq n}$ is assumed to be independent and identically distributed (i.i.d.), with the same distribution as a generic random vector (T, U) .

The independence assumption may not be entirely true, since some of the policyholders may be in contact, and thus can be able to transmit a malware more easily to one another. However, if we assume the portfolio to be large enough and if the subscription policy has avoided to constitute significant clusters of connected policyholders, such phenomena should be marginal and thus can be neglected. Typically this independence assumption reflects the fact that, if n is much smaller than the size of the national population (or even

global population), the infection is more likely to come from outside the portfolio than from inside. On the other hand, assuming that $(T_j, U_j)_{1 \leq j \leq n}$ are identically distributed may not be true if the heterogeneity of the portfolio is too strong. If it is the case, the portfolio should be segmented into risk classes.

Additionally, some countermeasures against the attack can exist. For example, the policyholders will progressively be warned about the threat and update their security system, making them "immune" to the attack. This process is similar to vaccination in classical epidemiology. We will denote by C_j the random time at which the policyholder j gains immunity. We again assume that these variables $(C_j)_{1 \leq j \leq n}$ are i.i.d. (with same distribution as a generic variable C), and independent from $(T_j, U_j)_{1 \leq j \leq n}$. The variables C_j act like right-censoring variables in survival analysis, see for example Fleming and Harrington (2011), in the sense that there is a "competition" between two time variables: if protection of the policyholder j occurs before the infection by the virus, T_j will never be observed. Let us define $Y_j = \inf(T_j, C_j)$, that is the time at which the policyholder j is no longer susceptible to be affected, either because immunity has been acquired, or because contamination has occurred. Introducing $\delta_j = \mathbf{1}_{T_j \leq C_j}$, the chronology of the events for policyholder j is summarized in Table 1 below.

Variable	Signification for policyholder j
T_j	Infection time
U_j	Duration of the assistance required by the victim
$T_j + U_j$	Time at which the assistance to the victim stops
C_j	Time at which policyholder j becomes immune
$\delta_j = \mathbf{1}_{T_j \leq C_j}$	Indicates if the policyholder j managed to become immune before infection

Table 1: Summary of the variables involved in the chronology of the cyber event at a policyholder level.

In the following, we will assume that T , U and C are independent. As the three variables (T, U, C) are duration variables, it is natural to define their distributions using their hazard rate. For a continuous random variable Z , the corresponding hazard rate function, denoted λ_Z in the following, is defined as

$$\lambda_Z(t) = \lim_{dt \rightarrow 0^+} \frac{1}{dt} \mathbb{P}(Z \in [t, t + dt] | Z \geq t).$$

The functions λ_U and λ_C reflect the response of the insurance company to the event. In

Section 4, we describe standard choices to model λ_U and λ_C , depending on the type of response and the reactivity of the company. On the other hand, λ_T is crucial since it describes the strength of the event itself and its dynamic. We now discuss in the following section how epidemiological models may be used to determine this λ_T .

2.2 Modeling the contagion

The distribution of T describes the force of the contagion. As we already mentioned in the previous section, we assume that contagion does not come from inside the portfolio itself, but from the outside. This approximation seems reasonable if the size of the portfolio is negligible compared to the global population of potential victims.

Hence the spread of the attack is defined as a global level, and it is quite natural to look at compartmental epidemiological models to describe this evolution of the environment. A standard choice is to consider SIR models. SIR models are commonly used to describe the evolution of epidemic, and more recently for modeling the spread of a malware through a network. SIR stands for "Susceptible" (exposed individuals), "Infected" (number of individuals affected by the disease) and "Recovered" (number of cured people). In some cases, the number of "recovered" people can also include individuals who died from the infection, and therefore can not be considered as contagious anymore. If s_t denotes the number of susceptible individuals in the population at time t , i_t the number of infected ones at time t and r_t the number of recovered ones, the model is described by the following set of differential equations

$$\frac{ds_t}{dt} = -\beta s_t i_t \quad (2.1)$$

$$\frac{di_t}{dt} = \beta s_t i_t - \gamma i_t \quad (2.2)$$

$$\frac{dr_t}{dt} = \gamma i_t \quad (2.3)$$

where β is the contagion rate and γ the "recovery" rate. The total size of the population, denoted by N , is then constant, with $N = s_t + i_t + r_t$. Also, in the simple model (2.1)-(2.3), recovered people do not become susceptible again. Many extensions have been proposed in the literature, to take vaccination into account, or the effect of various types of treatments.

From an initial value $s_0 \approx N$, $r_0 = 0$ and $i_0 \approx 0$, the system evolves to an epidemic or to an extinction of the crisis depending on the value of the "basic reproduction number"

defined as $R_0 = N\beta/\gamma$. If $R_0 > 1$, an epidemic appears, otherwise i_t vanishes rapidly. The dynamics of the epidemic hence depends on three parameters:

- β describing the strength of the contagion;
- γ a "recovery" parameter, which describes the rate at which an infected entity stops contributing to spread the malware;
- N the total number of entities exposed to the infection by the malware.

In our applications, we assume that

$$\lambda_T(t) = \beta i_t, \quad (2.4)$$

where i_t comes from the SIR model (2.1)-(2.3). This reflects the fact that our portfolio is a sample of the global population, whose evolution can be considered as deterministic because of its large size.

Remark 2.1 *Let us recall that $S_T(t) = \mathbb{P}(T \geq t)$, the survival function of T , can be expressed as $S_T(t) = \exp\left(-\int_t^\infty \lambda_T(u)du\right)$. The limit value $S_T(\infty)$ is not necessarily 0. In the special case where T can be infinite with non-zero probability (which means that some people will never be infected no matter how long the episode lasts), $S_T(\infty) = \mathbb{P}(T = \infty)$. Hence we are sure that the hazard function defined by (2.4) always corresponds to a probability distribution (potentially of a variable taking infinite values). Such variables are regularly considered in the survival analysis literature when it comes to studying "cure models", see Othus et al. (2012).*

2.3 Measuring the impact on the portfolio

Based on the description of the attack at a policyholder level, the insurance company is interested by the aggregation of these risks. We will introduce notations for the quantity we will be focusing on in the following. Let us define (recalling $\delta_j = \mathbf{1}_{T_j \leq C_j}$)

$$\begin{aligned} \mathfrak{N}_t &= \sum_{j=1}^n \delta_j \mathbf{1}_{Y_j \leq t} = \sum_{j=1}^n \delta_j \mathbf{1}_{T_j \leq t}, \\ \mathfrak{R}_t &= \sum_{j=1}^n \delta_j \mathbf{1}_{Y_j + U_j \leq t} = \sum_{j=1}^n \delta_j \mathbf{1}_{T_j + U_j \leq t}, \\ \mathfrak{I}_t &= \mathfrak{N}_t - \mathfrak{R}_t. \end{aligned}$$

$\mathfrak{N}_t$ denotes the cumulative number of infected policyholders at time t , while $\mathfrak{R}_t$ is the number of infected who have recovered before time t , and $\mathfrak{I}_t$ is the number for which the crisis is still ongoing. By "recovered", we only mean that these policyholders do not require a short-term assistance, and hence do not contribute anymore to saturation of the short-term response of the insurer. The total recovery may be much longer: in Section 4 we consider a cyber crisis who lasts a few days, while business disruption may be much longer, see Low (2017).

Let us now focus on the cost of the cyber event for the insurance company. We consider in this work three types of cost functions that can be combined:

$$\begin{aligned}\mathfrak{c}_1 &= \mathcal{C} \sup_{t \geq 0} \mathfrak{N}_t = \mathcal{C} \lim_{t \rightarrow \infty} \mathfrak{N}_t, \\ \mathfrak{c}_2 &= \mathbf{1}_{\sup_t \mathfrak{I}_t \geq K}, \\ \mathfrak{c}_3 &= \int_0^{t_d} \phi\left(\frac{\mathfrak{I}_t}{n}\right) dt,\end{aligned}$$

where $\mathcal{C}$ and K are positive constants, t_d is the duration of the attack and ϕ is a positive function. The cost function $\mathfrak{c}_3$ could equivalently be written in terms of the numbers of infected $\mathfrak{I}_t$, the writing with the proportion of infected $\mathfrak{I}_t/n$ being more convenient for the formulation of convergence results in Section 3. These different types of costs functions correspond to the different situations we wish to model:

- the cost function $\mathfrak{c}_1$ corresponds to a fixed cost per policyholder. In this case, only the total number of victims inside the portfolio is important, no matter the dynamics of the crisis. To simplify the situation, we did not consider random claim values, but $\mathcal{C}$ should be understood as a mean value for these claims. An adaptation to random claim values is straightforward but requires an additional modeling of the severity distribution.
- the function $\mathfrak{c}_2$ is introduced to describe the limited capacity of the insurance company to respond to the incident. The capacity of the assistance teams of the insurer can be overloaded if the number of policyholders to be helped becomes too large. This incapacity of the insurance company to intervene appropriately in a short amount of time may induce additional losses (financial penalties, loss of reputation, but also increased damages for the policyholders, left alone with no assistance). For example, if an insurance company can only help K policyholders at the same time, $\mathfrak{c}_2 = 1$ if the response system collapses.

- the cost function $\mathbf{c}_3$ is another way to consider saturation of the response. Here, we consider that there is a different cost per unit of time depending on the current number of policyholders requiring assistance. Indeed, such a situation may require to mobilize additional resources, and hence an additional cost.

In health risk analysis, the question to evaluate the cost for a portfolio, based on the evolution of the current number of infected people (which is our $\mathfrak{I}_t$) has been considered for example by Lefèvre et al. (2017), in a different framework (a compartmental stochastic SIR model). This corresponds to the particular case where ϕ in $\mathbf{c}_3$ is the identity function.

To analyse the behavior of these random functions $\mathbf{c}_1$ to $\mathbf{c}_3$, we need to get the distribution of $\mathfrak{N}_t$ and $\mathfrak{I}_t$ (and, consequently, of $\mathfrak{R}_t$). Section 3 hence aims at deriving approximations of these distributions.

3 Approximation of the evolution of the portfolio through Gaussian processes

To analyse the result of a cyber scenario on a portfolio, we need to determine the distribution of the evolution through time of infected policyholders and of those who have already been assisted by the company. Once whole distributions of the variables described in the framework Section 2 have been set, it is theoretically possible to simulate the whole evolution. In the present section, we derive asymptotic convergence results for the processes describing the evolution of the attack. The Gaussian approximations we derive allows us to understand the impact of such an attack on a large portfolio, either by avoiding to rely on simulations, or by simplifying the simulation process through the use of Gaussian variables. They also allow us to derive confidence bounds.

In Section 3.1, we first study the most simple problem, that is the distribution of the total number of infected policyholders during the episode. In this case, no understanding of the dynamics is required, only the final state of the portfolio. Section 3.2 shows how a Gaussian approximation for the processes involved in the infection can be obtained. As a corollary, Section 3.3 focuses on the cost function $\mathbf{c}_3$ which considers the case where the response cost per unit of time is function of the current proportion of infected policyholders.

3.1 Limit behavior of $\mathfrak{N}_t$

From the Law of Large Numbers, the limit behavior of the cumulated number of infected policyholders $\mathfrak{N}_t$ is given by

$$n^{-1} \sup_{t \geq 0} \mathfrak{N}_t = n^{-1} \lim_{t \rightarrow \infty} \mathfrak{N}_t = n^{-1} \sum_{j=1}^n \delta_j = \int_0^\infty S_C(t) f_T(t) dt = \nu \text{ a.s.}, \quad (3.1)$$

using the fact that T and C are independent, and where $S_C(t) = \mathbb{P}(C \geq t)$ is the survival function of the vaccination variable C and $f_T(t) = -S'_T(t)$ is the density of T .

The final proportion of victims is then the result of a competition between the variable T and the vaccination process C . A fast and efficient response C to the attack will lead to a fast decreasing S_C . On the other hand, if the spread of the attack is fast, which corresponds to a density f_T whose mass is concentrated near 0 (where S_C is close to 1), the proportion will be close to the totality of the portfolio.

On the other hand, from the Central Limit Theorem, we can easily get the asymptotic distribution of $\mathfrak{c}_1$, since

$$n^{-1/2} \left\{ \sup_{t \geq 0} \mathfrak{N}_t - n\nu \right\} \Longrightarrow \mathcal{N}(0, \nu(1 - \nu))$$

in which $\Longrightarrow$ denotes the convergence in distribution, when n tends to ∞ .

3.2 Gaussian approximation for $\mathfrak{N}_t$, $\mathfrak{R}_t$ and $\mathfrak{J}_t$

To study cost functions like $\mathfrak{c}_2$ and $\mathfrak{c}_3$, we do not only need the distribution of the final number of victims, but the distribution of the whole stochastic process describing the evolution of the epidemic. This multivariate process is defined as $\mathfrak{Z} = (\mathfrak{Z}_t)_{t \geq 0}$, with

$$\mathfrak{Z}_t = \begin{pmatrix} \mathfrak{N}_t \\ \mathfrak{R}_t \end{pmatrix}.$$

Let us define v^{tr} the transposition of a vector v . The central scenario is the expectation of $\mathfrak{Z}_t$, which is nz_t where $z_t = (\nu(t), \rho(t))^{tr}$, with

$$\nu(t) = \int_0^t S_C(u) f_T(u) du, \quad \text{and} \quad \rho(t) = \int_0^t S_C(u) f_V(u) du,$$

introducing f_T the density of T and f_V the density of $T + U$ (that is $f_T * f_U$ where $*$ denotes the convolution product and f_U the density of U). We define also

$$\eta_{Y,U}(t, h) = \int_0^t S_C(u) \{S_U(u) - S_U(t + h)\} f_T(u) du,$$

where S_U is the survival function of U . The deviations of $(\mathfrak{Z}_t)_{t \geq 0}$ from the central scenario are asymptotically distributed according to the results of Proposition 3.1 below.

Proposition 3.1 *Let $\mathfrak{Z}^c = (\mathfrak{Z}_t^c)_{t \geq 0}$ denote the centered process defined by*

$$\mathfrak{Z}_t^c = \frac{\mathfrak{Z}_t - nz_t}{n^{1/2}}.$$

Then $(\mathfrak{Z}_t^c)_{t \geq 0}$ converges in distribution towards a 2-dimensional Gaussian process $(\mathcal{Z}_t)_{t \geq 0}$

$$\mathfrak{Z}^c \Longrightarrow_{n \rightarrow \infty} \mathcal{Z}$$

where $\mathcal{Z}$ is centered with covariance structure

$$\begin{aligned} \Sigma(t, h) &= E \left[(\mathcal{Z}_t^{(1)}, \mathcal{Z}_t^{(2)})^{tr} (\mathcal{Z}_{t+h}^{(1)}, \mathcal{Z}_{t+h}^{(2)}) \right] \\ &= \begin{pmatrix} \nu(t)(1 - \nu(t+h)) & \eta_{Y,U}(t, h) - \nu(t)\rho(t+h) \\ \rho(t)(1 - \nu(t+h)) & \rho(t)(1 - \rho(t+h)) \end{pmatrix}, \end{aligned}$$

Proof. Let us introduce the class of functions

$$\mathcal{F} = \{ (t, c, u) \rightarrow \mathbf{1}_{t \leq c} (\mathbf{1}_{t \leq x}, \mathbf{1}_{t+u \leq x})^{tr} : x \in [0, \infty) \}. \quad (3.2)$$

Then $\mathfrak{Z}_t/n$ can be seen as the empirical measure applied to this class of functions. Hence, Proposition 3.1 is a straightforward application of Theorem 19.5 in van der Vaart (1998). Moreover, the process $\mathcal{Z}$ can be seen as a Gaussian process on $\mathcal{F}$, and is a $\mathbb{P}_C$ -Brownian bridge (see definition in Section 6.3), where $\mathbb{P}_C$ is a subprobability measure defined in section 6.3 (and related to the distribution of the response C). ■

The result of Proposition 3.1 has, at least, two implications: it shows the convergence at rate $n^{1/2}$ of our stochastic process towards a central scenario, and it provides an approximation for the deviations with respect to this scenario via the process $\mathcal{Z}$. In the simulation study of Section 4, the processes $(\mathfrak{Z}_t)_{t \geq 0}$ is easy to simulate directly (without relying on this approximation). Nevertheless, we explain how to simulate $\mathfrak{Z}$ in the general case in Section 6.3.

The asymptotic distribution of $(\mathfrak{Z}_t)_{t \geq 0}$ can be easily obtained as a corollary.

Corollary 3.2 *Let $\mathcal{Z}^\mathfrak{J} = (\mathcal{Z}_t^\mathfrak{J})_{t \geq 0}$ denote the centered Gaussian process with covariance structure*

$$\sigma_\mathfrak{J}(t, t+h) = \iota_1(t) - \iota_2(t)\iota_2(t+h),$$

where

$$\iota_1(t) = \nu(t) - \eta_{Y,U}(t, h), \quad \text{and} \quad \iota_2(t) = \nu(t) - \rho(t).$$

Then the centered process $\frac{(\mathfrak{J}_t - n\iota_2(t))}{n^{1/2}}$ converges in distribution towards $\mathcal{Z}^{\mathfrak{J}}$

$$t \rightarrow \frac{(\mathfrak{J}_t - n\iota_2(t))}{n^{1/2}} \Longrightarrow_{n \rightarrow \infty} \mathcal{Z}^{\mathfrak{J}}.$$

In particular, this means that the central scenario for the evolution of $(\mathfrak{J}_t)_{t \geq 0}$ is $t \rightarrow n\iota_2(t)$, with Gaussian errors around this trend. That is for a large portfolio, the proportion of infected policyholders is closed to $\iota_2(t) = \int_0^t S_C(u)(f_T(u) - f_V(u))du$. In this decomposition, the difference of the density of the infection times and the recovery times, is weighted by the survival function of the security implementation : the faster the vaccination, the smaller the proportion of infected policyholders.

Corollary 3.2 gives also an approximation of the whole distribution of $(\mathfrak{J}_t)_{t \geq 0}$, hence an approximation of $\mathbb{P}(\mathfrak{c}_2 = 1) = \mathbb{P}(\sup_{t \geq 0} \mathfrak{J}_t \geq K)$. Distribution-free deviation bounds can also be derived in Proposition 3.3.

Proposition 3.3 *For all $x \geq 0$,*

$$\mathbb{P} \left(n^{-1/2} \sup_{t \geq 0} |\mathfrak{J}_t - n\iota_2(t)| \geq x \right) \leq 2.5 \exp(-2x^2 + \mathcal{C}x),$$

for some absolute constant $\mathcal{C}$.

The proof is postponed to the appendix section, see Section 6.1. The constant $\mathcal{C}$ is hard to track, nevertheless its contribution becomes negligible when x is large.

3.3 Asymptotic distribution for $\mathfrak{c}_3$

The asymptotic distribution of $\mathfrak{c}_3$ is a consequence of Proposition 3.1. Indeed, let Φ denote the application

$$\Phi : \iota_2 \rightarrow \int_0^{t_d} \phi(\iota_2(t))dt. \quad (3.3)$$

The application Φ is Hadamard differentiable (see van der Vaart (1998) and section 6.2), and the asymptotic distribution of $\mathfrak{c}_3$ is then a consequence of the functional delta method (see Theorem 20.8 in van der Vaart (1998)). These arguments are detailed in Section 6.2 of the appendix, leading to the following Proposition 3.4.

Proposition 3.4 Assume that the function $\phi : [0, 1] \rightarrow \mathbb{R}$ is continuously differentiable with $\|\phi'\|_\infty < \infty$. Then

$$n^{1/2} \left\{ \mathfrak{c}_3 - \int_0^{t_d} \phi(\iota_2(t)) dt \right\} \Longrightarrow \mathcal{N}(0, \sigma^2),$$

where

$$\sigma^2 = \int_0^{t_d} \int_0^{t_d} \phi'(\iota_2(t)) \phi'(\iota_2(u)) \{ \iota_1(t \wedge u) - \iota_2(t \wedge u) \iota_1(t \vee u) \} dt du,$$

introducing the notation $t \wedge u = \min(t, u)$ and $t \vee u = \max(t, u)$.

3.4 Explicit computations for constant intensities

We provide here some explicit computations in the particular case of constant intensities λ_T , λ_U and λ_C for the infection time, the recover period and the immunization time. Then (for $\lambda_T \neq \lambda_U$)

$$S_C(t) = e^{-\lambda_C t}, \quad f_T(t) = \lambda_T e^{-\lambda_T t} \quad \text{and} \quad \nu(t) = \frac{\lambda_T}{\lambda_T + \lambda_C} (1 - e^{-(\lambda_T + \lambda_C)t}).$$

By convolution, the density of the random variable $V = T + U$ is given by

$$f_V(t) = \frac{\lambda_T \lambda_U}{\lambda_T - \lambda_U} (e^{-\lambda_U t} - e^{-\lambda_T t}), \quad \rho(t) = \frac{\lambda_T \lambda_U}{\lambda_T - \lambda_U} \left(\frac{1 - e^{-(\lambda_U + \lambda_C)t}}{\lambda_U + \lambda_C} - \frac{1 - e^{-(\lambda_T + \lambda_C)t}}{\lambda_T + \lambda_C} \right).$$

$$\eta_{Y,U}(t, h) = \lambda_T \left(\frac{1 - e^{-(\lambda_T + \lambda_C + \lambda_U)t}}{\lambda_T + \lambda_C + \lambda_U} - e^{-\lambda_U(t+h)} \frac{1 - e^{-(\lambda_T + \lambda_C)t}}{\lambda_T + \lambda_C} \right).$$

Thus, for a large portfolio, the proportion of infected policyholders is approximated by

$$i_2(t) = \frac{\lambda_T}{\lambda_T - \lambda_U} \left(\frac{\lambda_T}{\lambda_T + \lambda_C} (1 - e^{-(\lambda_T + \lambda_C)t}) - \frac{\lambda_U}{\lambda_U + \lambda_C} (1 - e^{-(\lambda_U + \lambda_C)t}) \right). \quad (3.4)$$

The proportion of infected policyholders increases up to the peak reached at time $\frac{\log(\lambda_T) - \log(\lambda_U)}{\lambda_T - \lambda_U}$, and then decreases. Remark that the immunization intensity λ_C does not impact this peak time, but it impacts the peak value $i_2(\frac{\log(\lambda_T) - \log(\lambda_U)}{\lambda_T - \lambda_U})$ which is decreasing in λ_C : the faster the immunization, the smaller the proportion of infected policyholders.

Remark: For $\lambda_T = \lambda_U$, taking the limit for $\lambda_U \rightarrow \lambda_T$ in the previous formula, or a direct computation leads to a $\Gamma(2, \frac{1}{\lambda_T})$ distribution for $V = T + U$ ($f_V(t) = \lambda_T^2 t e^{-\lambda_T t}$) and

$$i_2(t) = \frac{\lambda_T^2}{\lambda_T + \lambda_C} t e^{-(\lambda_T + \lambda_C)t} + \frac{(\lambda_T \lambda_C)}{(\lambda_T + \lambda_C)^2} (1 - e^{-(\lambda_T + \lambda_C)t}). \quad (3.5)$$

Then the proportion of infected policyholders increases up to the peak

$$\left(\frac{\lambda_T}{\lambda_T + \lambda_C} \right)^2 \left(\frac{\lambda_C}{\lambda_T} + e^{-\frac{(\lambda_T + \lambda_C)}{\lambda_T}} \right)$$

reached at time $\frac{1}{\lambda_T}$.

4 Simulation procedure

In this section, we illustrate how the simulation of a Wannacry-type scenario can be conducted, and how one can measure the efficiency of the response. The calibration of the Wannacry scenario is explained in Section 4.1. The different type of responses we consider are described in Section 4.2. The simulation results are gathered in Section 4.3.

4.1 Calibration of a Wannacry-type scenario

In this section, we discuss the parameters used for the contagion in our simulation setting. Rather than taking some arbitrary values for the parameter of the SIR model, we try to mimic an emblematic cyber-crisis episode.

Wannacry (Mohurle and Patil (2017)) is a famous ransomware global attack that stroke in May 2017. The hackers used the EternalBlue exploit, see Kao and Hsiao (2018). The ransomware propagated via Microsoft Windows users who did not patched their system against this vulnerability. More than 200 000 computers were affected by this attack, leading to massive immediate losses and business interruptions.

Using a SIR model to replicate the attack is a difficult task, since few available data support the calibration of the parameters. The SIR model depends on two parameters (β, γ) , the initial number of infected i_0 , and the size of the exposed population N .

In the description of classical epidemics, the parameter γ is linked to the time after which infected people move to the "recovered" category. In our case, it should not be understood as the time of full recovery (business interruption may last a long period), but as the time at which the infected entity stops being "contagious". In case of a Wannacry-type ransomware, one may consider that contagion is stopped relatively fast by introducing containment measures. That is why, in the following, we take $1/\gamma = 1$ day. On the other hand, alternative choices can be made for γ reflecting the fastness of the response. For example, in case of a dormant infection, $1/\gamma$ may be larger to reflect the time to detect the problem.

The initial number of infected entities i_0 is not a real problem, since it is supposed to be taken small, and then grow once the dynamics of the epidemic has fully started.

On the other hand, taking reasonable values for β and N is a much harder task. Indeed, there is no precise public data on the intensity of the contagion. On the other hand, the total exposure to risk, N , is unknown. A possible (disputable) way to proceed would be to consider that these susceptibles are all Windows users who did not update

their system. Even if the exact number of Windows users were easy to obtain, evaluating the ones that did not update their system is almost impossible. Rather than looking into this direction, we tried to fit this number using other characteristics of the SIR model and of the publicly available data on the Wannacry episode.

Introducing $i_{\max} = \sup_t i_t$ and recalling the notation $R_0 = N\beta/\gamma$ (the "Basic Reproduction Ratio"), we use the two following relationship for a SIR model,

$$i_{\max} = S_0 \left(1 - \frac{1 + R_0}{R_0} \right), \quad (4.1)$$

$$r_{\infty} = \lim_{t \rightarrow \infty} r_t = S_0 \left(1 - \exp \left(-\frac{\beta}{\gamma} r_{\infty} \right) \right). \quad (4.2)$$

From the knowledge of $i_{\max}$ and r_{∞} , we are able to find back S_0 and β , under the assumption that the SIR dynamic holds.

The quantity r_{∞} is the total number of infected, since, in the SIR model, every infected ends up as recovered. According to Chen and Bridges (2017), the Wannacry made approximatively 300,000 victims, although the exact number may be difficult to properly evaluate. On the other hand, determination of $i_{\max}$ requires to have a knowledge of the real-time evolution of the epidemic. Few data are (at least publicly) available to track the function $t \rightarrow i_t$ directly, except when it comes to the payment of ransoms (who is much less than the total estimated losses, see for example Field (2018)). Indeed, the ransomware was asking victims for bitcoins on three distinct addresses, see Willman (2017). The transactions are public, and allow to see at which time and date the ransoms have been paid. Of course, all victims did not pay a ransom (the total number of paid ransoms between 12th and 21th May 2017 is 320), but we may consider that the evolution of the number of payments reflects the kinetics of the epidemic. Let P_t denote the number of ransoms paid on day t , a rough assumption consists in assuming that $P_t = \alpha i_t$, where α is a fix proportion. The ratio between the total number of ransoms and the total number of victims leads to $\alpha = 937.5$. Moreover, the supremum of $t \rightarrow P_t$ is achieved on 15th May with 93 paid ransoms, leading to $i_{\max} = 87,188$. These values lead to the set of parameters and characteristics in the corresponding SIR model shown in Table 2.

The corresponding estimated function $t \rightarrow i_t$ during the ten first days of the Wannacry crisis is shown in Figure 1. Let us observe that, if we compare to the modeling of epidemics in human epidemiology, the Basic Reproduction Ratio R_0 we obtain is very close to 1. This quantity is commonly used in epidemiology to describe the contagiousness of a disease: an epidemic can start only if $R_0 > 1$ (which is the case here, but one is very close to the

	Value
β	2.556×10^{-7}
γ	1
N	4064279
R_0	1.04
$i_{\max}$	87188
r_{∞}	300,000

Table 2: Parameters and main characteristics for a SIR model calibrated from the Wannacry ransomware attack.

limit). Moreover, the larger R_0 , the wider the disease spreads, see for example Lefèvre and Picard (1996) for more details.

Figure 1: Corresponding function $t \rightarrow i_t$ estimated for the first ten days of Wannacry using the parameters of Table 2.

Again, we do not claim the SIR model of Table 2 to be accurate to describe the spread of the Wannacry epidemic, the calibration of the parameters and the underlying assumptions are too rough for that. Our objective is only to design a setting for our simulation that is reasonable according to past events like Wannacry, and the discussion above only aims at explaining in which sense we consider it "reasonable".

4.2 Behavior of the response

The response to the attack is linked to the variables U and C . In this section, we describe the set of distributions we use to describe different types of behaviors.

The variable U describes the response after infection, that is how long it takes for the infected policyholder to recover from the attack. What we mean by "recovery" may depend on the practical situation. According to studies like Low (2017), a full recovery after an attack may be a matter of weeks, months, even years for some victims. Even immediate business interruptions may be quite long. Since we are interested in the immediate response to a cyber crisis, the mean values of U we consider will be of the same magnitude as the length of the epidemic (which is around 10 days in case of a Wannacry-type scenario). This reflects the fact that we are interested in evaluating the potential saturation of the assistance that the insurance company brings to its policyholders. Hence, U has to be understood, in our case, as the time devoted to short-term measures in the first instants of the infection. We consider exponential distributions for U , and different values of the parameter to distinguish scenarios with fast or slow response.

The variable C describes the ability to react to the crisis. It reflects the fastness to identify the incident and to communicate on countermeasures that may prevent the spread of the attack through the portfolio. It also reflects the behavior of the policyholders, in their way to be immediately receptive or not to the alerts.

We consider three scenarii regarding the response variable C :

- a translated exponential distribution, $\lambda_C^{(1)}(t) = c_1 \mathbf{1}_{t \geq \tau_1}$. This means that, once the response has begun (at time τ_1), the proportion of policyholders per time who update their security system is constant (equal to c_1) through time;
- a Pareto-type distribution, $\lambda_C^{(2)}(t) = c_2(t - \tau_2 + 1/2)^{-\alpha_2} \mathbf{1}_{t \geq \tau_2}$, for $\alpha_2 > 0$. This corresponds to a situation where the vigilance of the policyholders decreases through time: the more careful ones perform update short after the date of response τ_2 , while the ones who did not instantaneously perform this update are more likely to ignore the threat;
- a Weibull-type distribution, $\lambda_C^{(3)}(t) = c_3(t - \tau_3)^{\alpha_3} \mathbf{1}_{t \geq \tau_3}$, for $\alpha_3 > 0$. This corresponds to a progressive attention devoted to this threat among policyholders.

In each case, the parameters $(\tau_j)_{1 \leq j \leq 3}$ represents the reactivity of the response.

4.3 Simulation results

We consider two portfolios of respectively $n = 5000$ and $n = 10,000$ exposed policies. For each portfolio, we perform 10,000 simulations of the impact of a cyber epidemic with the same attack intensity as Wannacry. For each type of response, we consider three delays of reaction: a fast response ($\tau_j = 3$ days after the start of the event), a medium response ($\tau_j = 5$ days), and a slow response ($\tau_j = 7$ days). The values of the parameters of the three types of responses described above are taken so that $E[C_j - \tau_j | C_j \geq \tau_j] = 1$, that is the response only differs by the shape of its hazard function (thus making comparisons more legitimate).

For each replication, we focus on the three types of cost functions $\mathbf{c}_j$, $j = 1, \dots, 3$ described in Section 2.3.

4.3.1 Cost function $\mathbf{c}_1$.

The cost function $\mathbf{c}_1$ is proportional to the total number of policyholders that have been affected. Summary statistics are shown in Table 3 below. One can observe that a fast response can reduce up to 63% the number of affected policyholders, while this reduction is only around 23% for a slow response. The magnitude of the reduction is of the same order between the different types of distributions for C . Nevertheless, one can observe that the Pareto reaction gives globally better result.

Type of reaction	Mean	Standard deviation	Median	Min	Max
No reaction	737.63	26.20	737	625	860
Slow Exponential	596.50	23.73	596	486	694
Slow Pareto	575.21	23.37	575	486	666
Slow Weibull	595.54	23.62	596	503	693
Medium Exponential	453.06	21.06	453	382	536
Medium Pareto	427.33	20.49	427	357	505
Medium Weibull	448.74	20.95	449	377	527
Fast Exponential	301.54	17.22	301	240	380
Fast Pareto	274.41	16.44	274	215	348
Fast Weibull	296.22	17.16	296	231	373

Table 3: Summary statistics for the total number of victims from 10,000 simulations ($n = 10,000$) depending on the reaction.

4.3.2 Cost function c_2 .

The cost function c_2 focuses on the maximum number of policyholders requiring immediate assistance, that is $\sup_t \mathfrak{I}_t$. Two typical simulated trajectories of $(\mathfrak{I}_t)_{t \geq 0}$ are shown in Figure 2. Some empirical statistics on $\sup_{t \geq 0} \mathfrak{I}_t$ are shown in Tables 4 and 5 below. The confidence intervals computed in these tables are bilateral: after ordering the 10,000 values for $\sup_{t \geq 0} \mathfrak{I}_t$, the left bound of the interval is the 250-th value, and the right-hand side is the 9750-th (in case of a 95% confidence interval). In other words, the upper bound K we obtain for the 95% confidence interval satisfies approximatively $\mathbb{P}(\sup_{t \geq 0} \mathfrak{I}_t \geq K) = 0.025$.

Figure 2: Two simulated trajectories of $t \rightarrow \mathfrak{I}_t$ (left-hand side: size of the portfolio $n = 5000$, right-hand side $n = 10,000$.)

Type of reaction	Mean	Standard deviation	95% confidence interval
No reaction	114.7122	8.91404	(98,133)
Slow Exponential	108.8625	9.241057	(92,128)
Slow Pareto	109.8391	9.269941	(93,129)
Slow Weibull	110.0568	9.213971	(93,129)
Medium Exponential	98.1915	9.220256	(81,117)
Medium Pareto	99.9944	9.378708	(82,119)
Medium Weibull	100.5861	9.292267	(83,119)
Fast Exponential	78.3415	8.359442	(63,95)
Fast Pareto	80.9131	8.712046	(64,98)
Fast Weibull	82.4775	8.667168	(66,100)

Table 4: Summary statistics on $\sup_{t \geq 0} \mathfrak{J}_t$, $n = 5000$.

Type of reaction	Mean	Standard deviation	95% confidence interval
No reaction	222.469	12.78206	(198,248)
Slow Exponential	212.6108	13.34208	(187,240)
Slow Pareto	214.622	13.38092	(189,242)
Slow Weibull	214.7712	13.2952	(190,242)
Medium Exponential	192.366	13.09089	(167,219)
Medium Pareto	196.3496	13.29188	(171,223)
Medium Weibull	197.0575	13.19889	(172,224)
Fast Exponential	153.595	11.87992	(131,177)
Fast Pareto	159.5131	12.33206	(136,184)
Fast Weibull	162.0158	12.23092	(138,187)

Table 5: Summary statistics on $\sup_{t \geq 0} \mathfrak{J}_t$, $n = 10,000$.

Additionally, we show the histograms for the variable $\sup_t \mathfrak{I}_t$ in the different settings for $n = 10,000$, in Figure 3 and 4. We see that all three types of responses lead to a similar impact on $\sup_t \mathfrak{I}_t$ (which is not entirely surprising since the expectation of these three distributions has been taken identical). Some differences in terms of variance still exist. The main parameter seems to be the time of response. A slow response will hardly diminish the burden of the assistance teams (reduction of around 4% only), while a fast response in 3 days significantly reduces (up to 30%) the magnitude of $\sup_t \mathfrak{I}_t$.

Figure 3: Histogram for $\sup_t \mathfrak{I}_t$ (10,000 simulations, size of the portfolio $n = 10,000$) in case of absence of response to the attack.

Figure 4: Histogram for $\sup_t \mathcal{J}_t$ (10,000 simulations, size of the portfolio $n = 10,000$) for different types of responses and delays.

4.3.3 Cost function $\mathbf{c}_3$.

As already mentioned, a particular situation consists in considering $\phi(u) = \alpha u$ for some $\alpha > 0$ in the definition of $\mathbf{c}_3$. This corresponds to the classical situation where a policyholder is assumed to generate the same cost per time of assistance, no matter the number of others requiring assistance at the same time. The two types of functions ϕ we study in the following consists of introducing a capacity of response K . As long as the number of policyholder to assist at a given time stays under K , the cost per policyholder is the same. After reaching this capacity of response K , additional resources have to be mobilized, and the cost per policyholder becomes higher. More precisely, the two functions ϕ_l and ϕ_{exp} that we consider are

$$\begin{aligned}\phi_l(u) &= n \{u + (1 + a)(u - K/n)\mathbf{1}_{nu > K}\}, \\ \phi_{\text{exp}}(u) &= nu + \exp(a(nu - K))\mathbf{1}_{nu > K}.\end{aligned}$$

Those functions, written in terms of u the proportion of infected, have the equivalent formulation in terms of number of infected $i = nu$

$$\varphi_l(i) = \phi_l(nu) = i + (1 + a)(i - K)\mathbf{1}_{i > K}, \quad \varphi_{\text{exp}}(i) = \phi_{\text{exp}}(nu) = i + \exp(a(i - K))\mathbf{1}_{i > K}.$$

We consider two values for the slope a , $a = 0.3$ and $a = 0.5$. The influence of the slope for both corresponding cost functions can be seen in Table 6 and 7 below, where we took $K = 100$.

In Figures 5 and 6, we show how the threshold K has impact on the global cost in the different situations we considered. Hence, the potential benefits of increasing or not the capacity of response K can be quantified.

Type of reaction	Mean	Standard deviation	95% confidence interval
No reaction, $a = 0.3$	2478.938	174.0935	(2141.704,2823.295)
No reaction, $a = 0.5$	2614.677	188.9222	(2248.99,2988.31)
Slow Exponential, $a = 0.3$	2173.087	171.6736	(1839.464,2512.994)
Slow Exponential, $a = 0.5$	2282.245	186.2669	(1920.66,2651.64)
Slow Pareto, $a = 0.3$	2128.768	169.6647	(1800.44,2468.661)
Slow Pareto, $a = 0.5$	2234.205	183.9841	(1878.74,2602.65)
Slow Weibull, $a = 0.3$	2196.921	171.721	(1867.764,2536.462)
Slow Weibull, $a = 0.5$	2308.166	186.3041	(1950.745,2676.685)
Medium Exponential, $a = 0.3$	1613.999	146.8447	(1333.837,1907.814)
Medium Exponential, $a = 0.5$	1678.619	158.3834	(1376.475,1995.25)
Medium Pareto, $a = 0.3$	1540.346	141.0674	(1269.801,1821.101)
Medium Pareto, $a = 0.5$	1600.994	151.862	(1309.02,1903.835)
Medium Weibull, $a = 0.3$	1626.793	145.6381	(1347.454,1917.305)
Medium Weibull, $a = 0.5$	1693.389	156.9928	(1392.78,2005.565)
Fast Exponential, $a = 0.3$	977.4292	106.331	(776.874,1194.924)
Fast Exponential, $a = 0.5$	1000.037	113.2298	(787.26,1231.595)
Fast Pareto, $a = 0.3$	897.8224	99.27368	(711.538,1096.531)
Fast Pareto, $a = 0.5$	918.1532	105.3581	(721.52,1128.995)
Fast Weibull, $a = 0.3$	985.268	106.0769	(787.357,1199.369)
Fast Weibull, $a = 0.5$	1010.295	112.923	(800.205,1238.93)

Table 6: Summary statistics on the global loss using ϕ_l in function $\mathbf{c}_3$, $n = 10,000$, $K = 100$.

Type of reaction	Mean	Standard deviation	95% confidence interval
No reaction, $a = 0.3$	1840.934	106.9957	(1634.27,2051.824)
No reaction, $a = 0.5$	2041.347	136.2028	(1781.106,2313.277)
Slow Exponential, $a = 0.3$	1657.649	104.4508	(1454.838,1866.687)
Slow Exponential, $a = 0.5$	1810.415	130.3399	(1560.442,2072.158])
Slow Pareto, $a = 0.3$	1631.331	103.7804	(1431.185,1838.067)
Slow Pareto, $a = 0.5$	1779.667	129.1652	(1532.008,2041.084)
Slow Weibull, $a = 0.3$	1672.122	104.7121	(1471.087,1878.891)
Slow Weibull, $a = 0.5$	1829.272	130.9318	[1580.514,2090.615]
Medium Exponential, $a = 0.3$	1307.649	92.90756	(1129.475,1493.088)
Medium Exponential, $a = 0.5$	1392.445	110.7826	(1182.142,1614.805)
Medium Pareto, $a = 0.3$	1253.033	90.74329	(1078.413,1434.561)
Medium Pareto, $a = 0.5$	1333.28	107.613	(1128.313,1550.503)
Medium Weibull, $a = 0.3$	1311.385	92.76599	(1132.492,1495.747)
Medium Weibull, $a = 0.5$	1400.128	110.8163	(1188.073,1622.126)
Fast Exponential, $a = 0.3$	870.2776	73.73921	(727.8004,1017.968)
Fast Exponential, $a = 0.5$	896.8764	82.40487	(739.458,1063.679)
Fast Pareto, $a = 0.3$	801.4578	70.61689	(666.1171,940.8641)
Fast Pareto, $a = 0.5$	825.6652	78.37562	(677.5672,981.6228)
Fast Weibull, $a = 0.3$	866.5024	73.82209	(725.2235,1013.799)
Fast Weibull, $a = 0.5$	896.6537	82.78573	(739.9383,1063.005)

Table 7: Summary statistics on the global loss using ϕ_{exp} in function $\mathfrak{c}_3$, $n = 10,000$, $K = 100$.

Figure 5: Evolution of the average loss depending on the threshold K in case there is no reaction. The black (resp. blue) curves correspond to functions of type ϕ_l (resp. ϕ_{exp}). The size of the portfolio is $n = 10,000$.

Figure 6: Evolution of the ratio of the average loss under a fast reaction and the average loss with no reaction depending on the threshold K . Blue corresponds to an exponential response, black to a Pareto response, red to a Weibull response. On the left-hand (resp. right-hand) side, using function ϕ_l (resp. ϕ_{exp}). The size of the portfolio is $n = 10,000$.

5 Conclusion

In this paper, we proposed a general framework to describe the impact of a cyber event on an insurance portfolio, allowing to design accumulation scenarii to assess the resilience against such types of massive attacks. The framework is flexible, since we distinguish between a global dynamic - here described through compartmental epidemiological models, but alternative models can easily be used - and the result at the portfolio level, which can be easily described through Gaussian approximation. We also gave a particular attention to different type of cost functions, which do not only take into account the total number of victims, but also the additional cost caused by a large number of policyholders requiring assistance simultaneously. Our model is also a tool to quantify the gain obtained through different strategies of reaction and/or prevention. Additional modifications could be done to enrich the model. In the present form, we considered, to simplify, a "perfect" reaction: once the risk has been identified, and means to act against it have been reported to the policyholders, they implement the countermeasures perfectly (after some delay, but then they are considered as fully protected). In practice, this may not be true, for example with large organizations that may not be able to identify and correct all their vulnerabilities. In which case, a reduction of the risk to be infected will be observed, but not necessarily a total cancellation of this risk. Moreover, a careful attention must be brought to the definition of realistic global scenarii to describe such a cyber epidemic. For now, there is a lack of public data to clearly identify precise timelines of massive cyber attacks. The example we took, inspired by the Wannacry crisis, should only be considered as a rough benchmark, due to the number of assumptions that have been done to try to mimic this episode.

6 Appendix

6.1 Proof of Proposition 3.3

The proof of Proposition 3.3 relies on Theorem 5 in Bitouzé et al. (1999). A version of this theorem adapted to our situation is summarized in Proposition 6.1 below. Before stating this result, we need to define the $L^2(\mathbb{P})$ -entropy with bracketing of a class of functions $\mathcal{F}$. An $L^2(\mathbb{P}) - \varepsilon$ -bracket is a set of functions $[l, g] = \{k : l \leq k \leq g\}$, where l and g are two functions such that $l \leq g$ $\mathbb{P}$ -almost surely, and $\int (l(\omega) - g(\omega))^2 d\mathbb{P}(\omega) \leq \varepsilon^2$.

The $L^2(\mathbb{P})$ –entropy with bracketing is the function $\varepsilon \rightarrow H_{\mathcal{F}}(\varepsilon) = \log(N_{[]}(\varepsilon, L^2(\mathbb{P}), \mathcal{F}))$ where $N_{[]}(\varepsilon, L^2(\mathbb{P}), \mathcal{F})$ is the smallest number of ε –brackets required to cover the class $\mathcal{F}$ (i.e. there exists a sequence of ε –brackets $[l_i, g_i]$ such that $\mathcal{F} \subset \bigcup_{i=1}^{N_{[]}(\varepsilon, L^2(\mathbb{P}), \mathcal{F})} [l_i, g_i]$). $L^2(\mathbb{P})$ –entropy with bracketing is a classical way to measure the complexity of a class of functions, see for example van der Vaart (1998) for more details.

Proposition 6.1 *Let $(Z_i)_{1 \leq i \leq n}$ denote an i.i.d. random vector. Let $\mathcal{F}$ denote a class of functions such that*

$$\forall f \in \mathcal{F}, 0 \leq f \leq M.$$

Let $H_{\mathcal{F}}$ denote the $L^2(\mathbb{P})$ –entropy with bracketing of $\mathcal{F}$, and assume that $H_{\mathcal{F}}(\varepsilon) \leq \gamma \varepsilon^{-1}$ for some $\gamma \in \mathbb{R}$. Then there exists an absolute constant C such that

$$\mathbb{P} \left(\sup_{f \in \mathcal{F}} \left| \frac{\sum_{i=1}^n f(Z_i) - E[f(Z_i)]}{n^{1/2}} \right| \geq \lambda \right) \leq 2.5 \exp(-2\lambda^2/M^2 + C\lambda).$$

To derive Proposition 3.3, we need to apply Proposition 6.1 to the class of functions

$$\mathcal{F} = \{(t, c, u) \rightarrow \mathbf{1}_{t \leq c} [\mathbf{1}_{t \leq x} - \mathbf{1}_{t+u \leq x}] : x \geq 0\}.$$

Consider $([l_i^{(1)}, g_i^{(1)}])_{1 \leq i \leq n_1}$ (resp. $([l_i^{(2)}, g_i^{(2)}])_{1 \leq i \leq n_2}$) a $L^2(\mathbb{P}) - (\varepsilon/2)$ –bracket of the class of functions $\mathcal{F}_1 = \{t \rightarrow \mathbf{1}_{t \leq x}\}$ (resp. $\mathcal{F}_2 = \{(t, u) \rightarrow \mathbf{1}_{t+u \leq x}\}$) where, from example 19.6 in van der Vaart (1998), $\max(n_1, n_2) \leq 4/\varepsilon$. For a given $x \geq 0$, there exists (i_1, i_2) such that

$$\begin{aligned} l_{i_1}^{(1)}(t) &\leq \mathbf{1}_{t \leq x} \leq g_{i_1}^{(1)}(t), \\ l_{i_2}^{(2)}(t) &\leq \mathbf{1}_{t+u \leq x} \leq g_{i_2}^{(2)}(t). \end{aligned}$$

Hence, define

$$\begin{aligned} l_{i_1, i_2}(t, c, u) &= \mathbf{1}_{t \leq c} (l_{i_1}^{(1)}(t) - g_{i_2}^{(2)}(t)), \\ g_{i_1, i_2}(t, c, u) &= \mathbf{1}_{t \leq c} (g_{i_1}^{(1)}(t) - l_{i_2}^{(2)}(t)). \end{aligned}$$

By construction, the sets of brackets $[l_{i_1, i_2}(t, c, u), g_{i_1, i_2}(t, c, u)]$, for all $i_1 \leq n_1$ and $i_2 \leq n_2$ forms an $L^2(\mathbb{P}) - \varepsilon$ –bracket of $\mathcal{F}$. Its size is less than $n_1 n_2$, which shows that

$$H_{\mathcal{F}}(\varepsilon) \leq \log(16/\varepsilon^2).$$

Hence Proposition 6.1 applies, and the result of Proposition 3.3 follows.

6.2 Proof of Proposition 3.4 using functional Delta-Method

Let $\mathbb{D}$ denote the set of bounded "cadlag" functions (see, for example, van der Vaart (1998) p.257, cadlag means right continuous functions whose limits from the left exist everywhere). A map $\Psi : \mathbb{D} \rightarrow \mathbb{R}$ is said Hadamard differentiable at $f \in \mathbb{D}$ if there exists $\Psi'_f : \mathbb{D} \rightarrow \mathbb{R}$ such that, for all $h \in \mathbb{D}$ and for any direction h_x (allowed to change with x) such that $\|h_x - h\|_\infty \rightarrow_{x \rightarrow 0} 0$,

$$\left| \frac{\Psi(f + xh_x) - \Psi(f)}{x} - \Psi'_f(h) \right| \rightarrow_{x \rightarrow 0} 0.$$

The function $\Phi : \iota_2 \rightarrow \int_0^{t_d} \phi(\iota_2(t))dt$ defined in (3.3) is a map from $\mathbb{D}$ to $\mathbb{R}$, and under the assumptions of Proposition 3.4 on ϕ , it is Hadamard differentiable at ι_2 . Indeed, from a first order Taylor-expansion,

$$\int_0^{t_d} \phi(\iota_2(t) + xh_x(t))dt - \int_0^{t_d} \phi(\iota_2(t))dt = x \int_0^{t_d} h_x(t)\phi'(g_h(t, x))dt,$$

where $g_h(t, x)$ is between $\iota_2(t)$ and $\iota_2(t) + xh_x(t)$. Hence, for all t , $h_x(t)\phi'(g_h(t, x)) \rightarrow_{x \rightarrow 0} h(t)\phi'(\iota_2(t))$. From Lebesgue's convergence theorem (since $\|\phi'\|_\infty < \infty$), we get that $\int_0^\tau h_x(t)\phi'(g_h(t, x))dt \rightarrow \int_0^\tau h(t)\phi'(\iota_2(t))dt$, leading to the existence of Φ'_{ι_2} defined as

$$\Phi'_{\iota_2} : h \in \mathbb{D} \rightarrow \int_0^{t_d} h(t)\phi'(\iota_2(t))dt.$$

From Corollary 3.2, $n^{1/2} \left\{ \frac{\mathfrak{J}}{n} - \iota_2 \right\} \implies \mathcal{Z}^\mathfrak{J}$. The functional Delta-method Theorem for Hadamard differentiable maps (see Theorem 20.8 in van der Vaart (1998)) then leads to

$$n^{1/2} \left\{ \Phi \left(\frac{\mathfrak{J}}{n} \right) - \Phi(\iota_2) \right\} \implies \Phi'_{\iota_2}(\mathcal{Z}^\mathfrak{J}),$$

leading to the result of Proposition 3.4.

6.3 $\mathbb{P}_C$ -Brownian motion on a set of functions

Consider a class of functions $\mathcal{G}$ taking values in $\mathbb{R}$. For a measure μ with finite mass, a μ -Brownian motion on $\mathcal{G}$ is defined as the centered Gaussian process $(W_\mu(\phi))_{\phi \in \mathcal{G}}$ with covariance structure

$$E[W_\mu(\phi)W_\mu^{tr}(\tilde{\phi})] = \int \phi(\mathbf{z})\tilde{\phi}^{tr}(\mathbf{z})d\mu(\mathbf{z}) = \langle \phi, \tilde{\phi}^{tr} \rangle_\mu.$$

In our case, we want to consider process on the class $\mathcal{F}$ defined in (3.2). However, we can observe that all functions in $\mathcal{F}$ are of the type $\mathbf{1}_{t \leq c}\phi(\mathbf{z})$, where ϕ belongs to the class

$$\mathcal{G} = \left\{ (t, u) \rightarrow (\mathbf{1}_{t \leq x}, \mathbf{1}_{t+u \leq x})^{tr} : x \in [0, \infty) \right\}.$$

Moreover, let us observe that (recalling that $\delta = \mathbf{1}_{T \leq C}$ and $Y = \inf(T, C)$)

$$E[\delta\phi(Y, U)] = E[S_C(T)\phi(T, U)] = \int S_C(t)\phi(t, u)d\mathbb{P}(t, u),$$

using the independence between C and (T, U) . Hence it is natural to define the limit processes as processes on $\mathcal{G}$, and to rely on the weighted measure $\mu = \mathbb{P}_C$, such that

$$\mathbb{P}_C(\phi) = \int S_C(t)\phi(t, u)d\mathbb{P}(t, u),$$

instead of $\mathbb{P}$.

Now, considering two classes of functions taking values in $\mathbb{R}$, say $\mathcal{G}_1$ and $\mathcal{G}_2$ containing the function $q_0 \equiv 1$, and $\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2$, define the process

$$\forall(\phi_1, \phi_2) \in \mathcal{G}_1 \times \mathcal{G}_2, B_{\mathbb{P}_C}(\phi_1, \phi_2) = \begin{pmatrix} W_{\mathbb{P}_C}(\phi_1) - \langle \phi_1, q_0 \rangle W_{\mathbb{P}_C}(q_0) \\ W_{\mathbb{P}_C}(\phi_2) - \langle \phi_2, q_0 \rangle W_{\mathbb{P}_C}(q_0) \end{pmatrix},$$

where $W_{\mathbb{P}_C}$ is a $\mathbb{P}_C$ -Brownian motion on $\mathcal{G}$. Each component of $B_{\mathbb{P}_C}$ is a $\mathbb{P}_C$ -Brownian bridge (see Khmaladze (2016)) on $\mathcal{G}_1$ or $\mathcal{G}_2$. The covariance structure of $B_{\mathbb{P}_C}$ is given by

$$E[B_{\mathbb{P}_C}(\phi_1, \phi_2)B_{\mathbb{P}_C}^{tr}(\tilde{\phi}_1, \tilde{\phi}_2)] = \begin{pmatrix} \mathbb{P}_C(\phi_1\tilde{\phi}_1) - \mathbb{P}_C(\phi_1)\mathbb{P}_C(\tilde{\phi}_1) & \mathbb{P}_C(\phi_1\tilde{\phi}_2) - \mathbb{P}_C(\phi_1)\mathbb{P}_C(\tilde{\phi}_2) \\ \mathbb{P}_C(\phi_2\tilde{\phi}_1) - \mathbb{P}_C(\phi_2)\mathbb{P}_C(\tilde{\phi}_1) & \mathbb{P}_C(\phi_2\tilde{\phi}_2) - \mathbb{P}_C(\phi_2)\mathbb{P}_C(\tilde{\phi}_2) \end{pmatrix},$$

Considering the classes of functions

$$\begin{aligned} \mathcal{G}_1 &= \{(t, u) \rightarrow \phi_{1,x}(t, u) = \mathbf{1}_{t \leq x} : x \in \mathbb{R}^+\}, \\ \mathcal{G}_2 &= \{(t, u) \rightarrow \phi_{2,x}(t, u) = \mathbf{1}_{t+u \leq x} : x \in \mathbb{R}^+\}. \end{aligned}$$

The process $(\mathcal{B}(x))_{x \geq 0}$ defined as $\mathcal{B}(x) = B_{\mathbb{P}_C}(\phi_{1,x}, \phi_{2,x})$ is a Gaussian process on $\mathbb{R}^+$, and it is easy to check that its covariance structure is the same as the process $\mathcal{Z}$ defined in Proposition 3.1. Hence, to simulate $\mathcal{Z}$, a simple way to proceed is to simulate $B_{\mathbb{P}_C}$, which itself can be deduced from the simulation of $\mathcal{W}(x) = (W_{\mathbb{P}_C}(\phi_{1,x}), W_{\mathbb{P}_C}(\phi_{2,x}))^{tr}$.

The simulation of $\mathcal{W}$ is relatively easy, since, for all $h > 0$, $\mathcal{W}(x+h) - \mathcal{W}(x)$ is independent from $(\mathcal{W}(x'))_{x' \leq x}$. Indeed, introducing the notation $\Delta_h \phi_{j,x} = \phi_{j,x+h} - \phi_{j,x}$ for $j = 1, 2$, we have

$$\mathcal{W}(x+h) - \mathcal{W}(x) \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \langle \Delta_h \phi_{1,x}, \Delta_h \phi_{1,x} \rangle & \langle \Delta_h \phi_{1,x}, \Delta_h \phi_{2,x} \rangle \\ \langle \Delta_h \phi_{2,x}, \Delta_h \phi_{1,x} \rangle & \langle \Delta_h \phi_{2,x}, \Delta_h \phi_{2,x} \rangle \end{pmatrix}\right).$$

The covariance matrix can be rewritten as

$$\begin{pmatrix} \int_x^{x+h} S_C(u)f_T(u)du & \int_x^{x+h} S_C(u)F_U(x+h-u)f_T(u)du \\ \int_x^{x+h} S_C(u)F_U(x+h-u)f_T(u)du & \int_x^{x+h} S_C(u)f_V(u)du \end{pmatrix}.$$

When h is small, this can be approximated by

$$h \times \begin{pmatrix} S_C(x)f_T(x) & , & S_C(x)F_U(h)f_T(x) \\ S_C(x)F_U(h)f_T(x) & , & S_C(x)f_V(x) \end{pmatrix}. \quad (6.1)$$

Hence, when it comes to simulating such a process, (6.1) can be used to simulate $\mathcal{W}$ through its increments. Then, $\mathcal{B}$ can be easily deduced.

Acknowledgement: *The authors acknowledge funding from the project Cyber Risk Insurance: actuarial modeling, Joint Research Initiative under the aegis of Risk Foundation, with partnership of AXA, AXA GRM, ENSAE and Sorbonne Université.*

References

- Artalejo, J. and Lopez-Herrero, M. (2014). Stochastic epidemic models: New behavioral indicators of the disease spreading. *Applied Mathematical Modelling*, 38(17-18):4371–4387.
- Bessy-Roland, Y., Boumezoued, A., and Hillairet, C. (2020). Multivariate hawkes Process for cyber insurance. *Preprint*.
- Bitouzé, D., Laurent, B., and Massart, P. (1999). A Dvoretzky-Kiefer-Wolfowitz type inequality for the Kaplan-Meier estimator. *Annales de l’I.H.P. Probabilités et statistiques*, 35(6):735–763.
- Camillo, M. (2017). Cyber risk and the changing role of insurance. *Journal of Cyber Policy*, 2(1):53–63.
- Chen, H. and Cox, S. H. (2009). An option-based operational risk management model for pandemics. *North American Actuarial Journal*, 13(1):54–76.
- Chen, Q. and Bridges, R. A. (2017). Automated behavioral analysis of malware: A case study of wannacry ransomware. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 454–460. IEEE.
- El Koufi, A., Adnani, J., Bennar, A., and Yousfi, N. (2019). Analysis of a stochastic sir model with vaccination and nonlinear incidence rate. *International Journal of Differential Equations*, 2019.

- Eling, M. and Loperfido, N. (2017). Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: Mathematics and Economics*, 75:126–136.
- Eling, M. and Schnell, W. (2016). What do we know about cyber risk and cyber risk insurance? *The Journal of Risk Finance*.
- Farkas, S., Lopez, O., and Thomas, M. (2019). Cyber claim analysis through Generalized Pareto Regression Trees with applications to insurance pricing and reserving. <https://hal.archives-ouvertes.fr/hal-02118080>.
- Field, M. (2018). Wannacry cyber attack cost the NHS £ 92m as 19,000 appointments cancelled. *The Telegraph*, page 2018.
- Fleming, T. R. and Harrington, D. P. (2011). *Counting processes and survival analysis*, volume 169. John Wiley & Sons.
- Forrest, S., Hofmeyr, S., and Edwards, B. (2016). Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity*, 2(1):3–14.
- Kao, D.-Y. and Hsiao, S.-C. (2018). The dynamic analysis of Wannacry ransomware. In *2018 20th International Conference on Advanced Communication Technology (ICACT)*, pages 159–166. IEEE.
- Kermack, W. and McKendrick, A. (1927). A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. A.*, 115:700–721.
- Khmaladze, E. (2016). Unitary transformations, empirical processes and distribution free testing. *Bernoulli*, 22(1):563–588.
- Lefèvre, C. and Picard, P. (1996). Collective epidemic models. *Mathematical Biosciences*, 134(1):51 – 70.
- Lefèvre, C., Picard, P., and Simon, M. (2017). Epidemic risk and insurance coverage. *Journal of Applied Probability*, 54(1):286–303.
- Low, P. (2017). Insuring against cyber-attacks. *Computer Fraud & Security*, 2017(4):18 – 20.
- McKendrick, A. G. (1925). Applications of mathematics to medical problems. *Proceedings of the Edinburgh Mathematical Society*, 44:98–130.

- Mohurle, S. and Patil, M. (2017). A brief study of Wannacry threat: Ransomware attack 2017. *International Journal of Advanced Research in Computer Science*, 8(5).
- Montagnon, P. (2019). A stochastic SIR model on a graph with epidemiological and population dynamics occurring over the same time scale. *Journal of mathematical biology*, 79(1):31–62.
- Othus, M., Barlogie, B., LeBlanc, M. L., and Crowley, J. J. (2012). Cure models as a useful statistical tool for analyzing survival. *Clinical Cancer Research*, 18(14):3731–3736.
- Runhuan Feng, J. G. (2011). Actuarial applications of epidemiological models. *North American Actuarial Journal*, 15(1):112–136.
- van der Vaart, A. W. (1998). *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- Willman, R. (2017). Wannacry outbreak data. <https://github.com/rwillmann/WannaCry-Outbreak-Data-12-May-2017---19-May-2017->.
- Zhang, X. and Wang, K. (2013). Stochastic SIR model with jumps. *Applied Mathematics Letters*, 26(8):867–874.