# Designing Organizations for Cyber Security Resilience

Marshall Kuypers[1] and Thomas Maillart[2]

[1] Qadium Inc., USA
[2] School of Economics and Management, University of Geneva, Switzerland

**Abstract.** Every day, security engineers cope with a flow of cyber security incidents. While most incidents require routine reactions, few require orders of magnitude more effort to investigate and resolve. It therefore remains unclear how security operation teams in organizations should tune their response to handle large flows of incidents and, at the same time, tame extreme events. Analyzing the statistical properties of 60,767 security incidents collected over more than six years at a large organization, we find that the distribution of resolution effort induced by security incidents is in general skewed, following a power law tail distribution with exponent $\approx 1.5$. However, this distribution of incident severity becomes less skewed over time, suggesting that the organization under scrutiny has managed to reduce the impact of large events. Thus, the organization could use the time saved to cope with a super-exponential increase of cyber security incidents, while keeping human resources stable overall. We also find that the flow of cyber security incidents triggers short-lasted effort spikes, which may be up to 2 orders of magnitude larger than median instant effort. Our results offer a first quantitative reference point on how cyber security incidents may affect organizations on the long term, and how organizations may adapt to efficiently absorb cybersecurity shocks.

## 1 Introduction

Despite the large stakes currently associated with cyber security, unavailability of incident data has impeded the monitoring, analysis, and forecast of cyber risks at the level of organizations. Information sharing initiatives are making progress,[3] but at this time, simple metrics about the rate and impact of cyber attacks have remained largely inaccessible to researchers and to the broader public. Some publicly available databases record cyber security incidents, but they are usually heavily biased towards incidents that require public disclosure and large-scale incidents at major organizations that are reported by the media [1]. Therefore, many open questions exist about the statistical properties of cyber security incidents and their implications for the organization of security response. Fortunately, cyber security incident data are collected for the purpose of auditing, compliance, and management [2]. These data can be leveraged to derive risk metrics, which in turn may lead to more informed security management and investments [3]. Here, we present a statistical analysis of 60,767 cyber security incidents spanning six years and two months that occurred at a large US organization.

[3] For example, see information sharing initiatives in [1, 2].

Like many other natural [4,5] and manmade disasters [6,7], cyber incidents involve extreme events [8,9,10]. Even if these events do not necessarily lead to existential damage, they may have unintended consequences. For instance, at a large Internet company, it was reported to the authors that the Heartbleed vulnerability[4] required a team of more than five expert security engineers over a whole week (i.e., approximately 250 hours) to safely deploy the necessary patches across the infrastructure, with obvious disruptive effects and delays on the normal workflow of cyber security incident management.[5] If even best equipped and security aware organizations are struggling with such extreme events, one may question the ability of average organizations facing similar challenges to quickly respond to critical incidents with enough manpower. Deciding how much in-house human expertise versus outsourcing [11], risk acceptance [3] and transfer (e.g., insurance) [12], is becoming an increasingly relevant question for the management of numbers of organizations.

In this paper, our objective is to bring quantitative insights on the relative weight of a few large cyber security incidents, compared to a multitude of small events. In the present data set, the 100 largest events (out of 60,767 events) account for 25% for the cyber security incident response effort (thereafter, the *effort* in man-hours). Calibrating the statistical properties of cyber security incidents requires extreme risk modelling, which has been used to address other outstanding cyber security challenges, such as heavy-tailed distributions of personal data breaches [8, 9, 10]. We find that cyber security incidents tend to become overall less extreme over time. In particular, we illustrate this latter point with lost or stolen devices (e.g., laptops), and we show how the organization studied here has addressed this challenge by implementing a full disk encryption (FDE) policy. We find that the reduction of extreme events may have helped the organization build capacity to absorb (resp. consider) a much larger amount of cyber security incidents with rather stable human resources. We also find that required effort is regularly punctuated by jumps of *excess effort*, which frequency of occurrence given amount of excess effort can be quantified. Finally, we show how the subsidiaries of the same large organization exhibit time dependence in terms of event frequency.

The rest of the paper is organized as follows. First, we review background research (Section 2). We then present the nature of cyber security incidents contained in our data set (Section 3), followed by a description of the standard statistical methods (Section 4). We then turn to results (Section 5), discussion (Section 6) and conclusion (Section 7).

## 2 Related Work

Most research in cyber security is motivated by established flaws or potential incidents, which may disrupt the normal operations of organizations active on the Internet. Typical research perspectives include documenting the origins and failure mechanisms of cyber

[4] See heartbleed.com for more information.

[5] Private conversation with a security engineer at a large payment processing company.

security incidents [13], as well as their often non-obvious economic and social consequences at people, organization, and country levels [14]. Popular security incidents include software vulnerabilities [15, 16], operation disruptions (in particular for critical infrastructures [13]), personal and sensitive data thefts [8, 10] as a result of Internet attacks [17], insider threats [18, 19], and human mistakes [20]. Most of these incidents carry their own uncertainties regarding probability of occurrence and severity, which most often remain hard to quantify since data are generally kept private by stakeholders. In some situations however, organizations are legally required to disclose publicly (e.g., personal data breaches), or given incentives to share security incidents with governmental agencies [21]. In many cases, public release has brought better understanding about the risks associated with these events [22, 23], such as robust statistical models of personal data breaches [8, 10], and predictive algorithms of software vulnerabilities [17, 20, 24, 25].

For other categories of cyber security incidents, the paucity of data has impeded progress of collective understanding of these events. As such, it has limited the development and widespread adoption of best practices, which would have the potential to improve collective benefits in comparable ways to vaccination in epidemiology [26]. For the time being, most organizations have no requirement and usually no incentive to disclose information about cyber security incidents, even though recording and documenting these incidents has long been recognized as a critical part of IT security [27, 28]. For instance, the US- CERT requires that certain information be recorded and reported when an incident occurs on a federal information system [21] with some clear guidelines.[6] Therefore, many organizations record data in incident management systems but may not fully leverage that information. Ahmad et al. [29] analyzed incident management systems at a financial institution and found that miscommunication and organizational barriers prevented incident data from being best used. Therefore, academic research is typically limited to either theoretical considerations or surveys about incidents and practice, instead of data-driven empirical analyses [2]. Nevertheless, some investigations have brought insights associated to cyber security incidents within organizations. In 2008, Condon et al. [30] published an analysis of security incidents at the University of Maryland. However, the data consists primarily of malware incidents, and only their frequency was studied but not their severity. Others have studied vulnerability disclosures [31], or even the failure of financial information systems [32], but work analyzing cyber incident data has remained scarce.

In this paper, we show how a longitudinal analysis of historical cyber security incidents can be used to obtain probability distributions of incident severity, which we measured here in *man-hours* of investigation, remediation and improvement. Kuypers and Paté-Cornell have used these probabilistic outputs as inputs to quantitative risk models to assess cyber risks in dollar terms by modelling the cost of incident investigation [33], and also reputation damage, business interruption, intellectual property loss, and other costs [12]. The generation of these probabilistic inputs is critical, given the heavy-tailed nature of some cyber incidents. Other cyber risk models have histori-

[6] See also https://www.us-cert.gov/incident-notification-guidelines

cally used expected values of losses instead of probabilistic inputs, probably because of data constraints [34, 35]. Models that use Monte Carlo simulations or other methods to probabilistically assess risk provide much more information about cyber risks [36].

## 3 Data

The data set contains 60,767 cyber security events, which have occurred at a large organization based in the United States of America over 6 years and 2 months, between November 2008 and January 2015. All recorded events were assigned a tracking number, the date of the incident, the number of systems impacted, the total number of hours of investigation, the suspected attackers, other details about the incident, and a resolution date for mainly the 38,147 first events. All incidents were collected through an incident tracking system, which is common in most mid-size and large organizations operating a security operations center (SOC). Information about events recorded in ticketing systems have been manually entered by security officers. A specificity of the data set used in this study are the records of *effort* in man-hours for each event, which generally includes time spent remediating the incident. For example, the costs of a malware infection investigation include the time that an analyst must spend to identify the malware, wipe the hard drive, and reload a data backup.

The data contain a wide range of incidents, including lost or stolen devices, denial of service attacks, network exercises, employee misuse, phishing attacks, malware infections, and unauthorized access by attackers. In case of cyber attacks, the perpetrators also represent nearly all categories of attackers, including insiders, hacktivists, criminals, and nation-states. Even though, each incident is unique in its nature, we can best describe their nature at the aggregate level with seven categories:

1. **Data spillages:** Incidents that possibly disclose information to unauthorized individuals. For example, an employee could forget to encrypt an email that contains social security numbers.
2. **Email incidents:** Any intrusion or attempted attack that originates through email is classified as an email incident. For example, criminals may try to extract a user's email credentials to use their account for sending spam (e.g., *phishing*), or attach malicious files to an email to infect a user's machine with malware.
3. **Lost or stolen devices:** Laptops, tablets, phones, and other hardware can be lost or stolen. These incidents typically require different levels of investigation depending on device type and encryption level.
4. **Tasks:** Incidents caused by network exercises, wide scale patching, or investigations (such as pulling log files) meant to aid an audit or an inquiry (e.g. pulling an employee's emails after allegations of harassment).
5. **Website incidents:** Any attack that exploits websites operated by the organization is classified as a website incident, including website defacements, SQL injections, and server compromises.
6. **Web browsing and USB incidents:** Malware that does not originate via email or through a website is categorized as a web browsing/USB incident. For example, a

user may inadvertently download malware while visiting a compromised website. Users may also spread malware via USB devices. This category excludes malware delivered via email.

7. **Other:** While other types of incidents occur, many are not frequent enough to fall into a specific category. For example, denial of service attacks and insider attacks occurred very rarely at the organization that we studied. False alarms and near misses were also reported. We consider this class of events as a category by itself, despite its heterogeneity. This, however, limits the conclusions that may be drawn from a specific analysis of this category.

These categories have been established in accordance to the perception of the co-author who cleaned the data. The method used for cleaning used a combination of (i) incident labels, (ii) regular expressions matching key words and phrases, (iii) manual review. The data contained incident categories, some of which were broken out as specific categories (e.g. DDoS attacks were explicitly labeled). Sub-tags were in use in some cases as well (e.g. stolen laptop). Further tags could be used to categorize more incidents (for example, some incidents had a *malware family* label, which helps categorize the incident). This incident label technique was useful to categorize roughly 80% of the records. Next, we looked for strings that we could match. For example, any incident containing *laptop* and *stolen* could easily be categorized as a lost device. Similar string matching categorized another 19%. The remaining 600 incidents were manually categorized by reading the incident description.

The categories obviously carry their load of subjectivity. Yet, they help bring a qualitative description of the events that have occurred over the span of the data set. Overall, the data offer a comprehensive view of the human resources deployed by the organization over a long period of time. They illustrate how this effort is distributed to tackle a large spectrum of cyber security incidents. The frequency of incidents is taken into account and their severity is measured as effort in man-hours. The time spent by security engineers (and other people who may be involved in subsequent crisis management and mitigation) may not be the only source of monetary costs of cyber incidents. Yet, we believe that the expenses of human resources represent a best-effort proxy of the effort required to overcome these incidents.

In addition, our data set contains information regarding different sites (respectively subsidiaries) by the large organization. We shall inspect the relevant statistical properties of each site, as well as their mutual dependence for the 15 sites for which we have more than 250 events recorded.

## 4 Method

Nowadays, most organizations face a continuous flow of cyber security incidents, including external attacks, insider attacks, maintenance (e.g., patch deployment) or stress testing tasks (e.g., vulnerability assessments). These incidents occur with frequency and severity that we measure here as the effort required to overcome the incident. Here, we

observe that resolution effort of some events are orders of magnitude larger than the median incident severity. The severity and the frequency of large events relative to average events is determinant and must be appropriately quantified [37].

### 4.1 Accounting for qualitatively different events

Our statistical method is crafted to account for the existence of 3 levels of event severity as observed in our data set: (i) small *routine* events, (ii) heavy-tail large events, and (iii) extreme outliers, which presumably stem from qualitatively different incidents [38]. For each type of incidents, we quantify the evolution of their frequency and severity over the six years by segmenting the data into 12 time periods of approximately 6 months, then analyzing them at the aggregate level.

### 4.2 Evolution of incident frequency and severity & return times

The community interested in extreme risks has long discussed the nature of tail risks and some popular methods have been developed in the past to identify the model that *best* fits heavy-tailed random samples [39, 40]. The usual point of debate is whether an *extreme* risk is actually extreme and bound to become more extreme over time (namely, with no statistical moment defined), or on the contrary, whether there is an upper limit of severity (see [10] for a study of personal data breaches, as a concrete example of an extreme, yet bounded risk). Overall, fitting extreme distributions is a challenge because statistics usually build on the law of large numbers, while extreme events are by definition rarely observed and, as such, sampled over large periods of time. A variety of tools, such as Extreme Value Theory (EVT), have been developed to assess the probability of an extreme event beyond observations [41]. These tools are particularly useful for the (re)insurance industry, because they bring robust forecasts about the maximum claim amount resulting from a large disaster, allow predictive models to be developed, and help compute competitive insurance premiums [42]. Here, we observe that event severity follows a power law tail distribution given by,[7]

$$P\left(S \geq s\right) \sim \frac{1}{x^{\mu}}, \tag{1}$$

where $s$ is the severity and $\mu$ its exponent, within the boundaries defined by the lower threshold, and for values smaller than the outliers. Based on empirical evidence evidence (see Figure 2), we assert the hypothesis that one cannot rule out that the tail data were drawn from the power law. Given a sample, we find the best fit using maximum likelihood estimator (MLE) [8, 39, 40]. Goodness-of-fit is obtained by performing Monte Carlo simulations of synthetic data sets with the same parameters and size (bootstrapping method), and by using the Kolmogorov-Smirnov (KS) test, which measures the maximum distance between the model and the generated distributions [43]. The *p-value* is obtained as the ratio of KS statistics for the synthetic data sets whose value exceeds the KS test for the real data set; therefore, the larger *p*, the more accurate the

[7] Note that $f\left(x\right) = 1/x^{\mu}$ implies $log\left(f\right) = -\mu \cdot log\left(x\right)$, therefore the linear relationship in double logarithmic scale.

model's description of the data. We chose a relatively conservative level ($p > 0.1$) as the rejection level for the null hypothesis [40]. The outliers are detected by removing extreme values, until $p > 0.1$ for the tail distribution of the main power law (i.e., until we can reject the null hypothesis). If more than a couple of outliers are present, we attempt to fit an outlier tail regime, with a power law model [10]. To account for the evolution of risks, we repeat the distribution fitting exercise over time periods of 6 months, which are sufficiently fine-grained to capture change with large enough sample for proper fitting. We additionally also measure the evolution of percentiles effort from fine-grained monthly bins.

### 4.3 Return times

We recognize that arrival of new events of sufficiently large severity (effort $S \geq 6$ man-hours) follows an exponential distribution. Therefore, the underlying process follows a memoryless Poisson process, and with arrival rate tested and varying at various thresholds.

### 4.4 Cyber security incident activity across sites

We first check the consistency of statistical properties across sites, and we then measure the activity dependence between sites. For this, we use Spearman rank dependence measurements [39]. These dependencies shall inform on the probability that a change of regime may affect multiple sites at once. Our approach demonstrates a prototype of portfolio management for cyber risks [39].

## 5 Results

Establishing ground for sound quantitative risk management requires to study the stability and the possible evolution of frequency and severity variables, as well as relevant dependences between these variables. With a data set of 60,767 cyber security incidents at a large organization, we uncover the statistical features associated with the management of cyber security incidents, as well as their evolution over the course of 6.2 years, a rather long period in the short history of the Internet. This section is organized as follows: (i) we show how the frequency of events has experienced a super-exponential increase and how the frequency of cyber security event type has evolved. (ii) We study the evolution severity $S$ (i.e., effort to resolve cyber security incidents in man-hours) over time, which find to evolve positively and become less heavy-tailed with time. For the sake of concreteness, we illustrate this evolution with a case study on the deployment of full-disk encryption (FDE) and how it has influenced positively risk management by reducing mitigation efforts associated with lost (respectively stolen) devices. (iii) Like for other risks (e.g., natural disasters), managers may want to measure return time of events above a severity threshold. We provide such return times for cyber security incidents for the large organization under scrutiny. (iv) The heavy-tailed nature of cyber security incident severity triggers large sudden jumps of effort above the baseline

work activity. We aim to understand how these jumps influence hourly workload. We find that baseline work is very stable, however, with sharp sudden jumps, which we call "excess effort". We find that this excess effort follows a power law tail distribution with exponent $\alpha_{pos} \approx 1.41$. (v) We finally investigate the dependence (i.e., correlations) of event frequency between 15 sites (respectively subsidiaries) with more than 250 events over 6.2 years at the large organization. Our results show how quantitative risk metrics (i.e, frequency x damage) could also help build cyber security risk portfolios across sub-organizations (within an organization) and across organizations, e.g., within a specific industry or country.

### 5.1 Evolution of cyber security incidents frequency

Figure 1B shows the evolution of event count (binned per month) per category of events (cumulative). While the first half of the 2008 until 2015 period was dominated by incident associated with malware from Web browsing and USB (purple), the second half period exhibits the rise of stolen devices (dark green) and incidents associated with websites (light green). These dynamics of event frequency show that the prominence of some categories tends to be replaced by emerging categories. Figure 1A shows the overall evolution of event counts (binned per week). The increase follows a super-exponential growth with finite-time singularity given by [44, 45, 46],

$$f(t) \sim (t - t_c)^{\nu}, \tag{2}$$

with $\nu = 0.52$ and $t_c = 6$. Finite-time singularities are the hallmark of positive feedback loops. In the context of cyber security incidents, finding such acceleration of event frequency is rather unexpected. This phenomenon can be attributed to capacity building at overall constant human resources: the reduction of large events frequency released available resources for the absorption of more – less severe – events.

The distribution of cyber incident resolution effort can be characterized by three regimes: (i) low effort incidents (less than 2 man-hours of work), (ii) a main tail section, well represented by a power law distribution, and in some cases, (iii) a transient outlier tail regime, also well described by a more extreme power law distribution with a smaller exponent. Figure 2B shows the Complementary Cumulative Distribution Function (CCDF) of all incidents that occurred during the 11th time period. The CCDF determines the probability that an event will cost more than a given amount. For instance, there is approximately 1% (resp. 0.1%) chance that an event involving more than 10 (resp. 100) hours of work will occur. The tail distribution shown in Figure 2B is best described by a main power law tail $P\left(C \geq c\right) \sim 1/c^{\mu}$ , with $\mu$ = 1.80 $\pm$ 0.06, for 0.5 $<$c $<$30, and an outlier tail regime with $\mu$ = 0.8 $\pm$ 0.5, for c $\geq$ 30. The outlier regime is of particular importance: while the main tail has an exponent $\mu > 1$ with its first statistical moment (i.e., average) defined, the outlier tail ($\mu < 1$) has no statistical moment defined. In other words, as more events get sampled, new even more extreme events are likely to appear, pushing the average towards larger values as a result of the outlier regime. Note that these outlier regimes are transitory over the 12 periods considered.

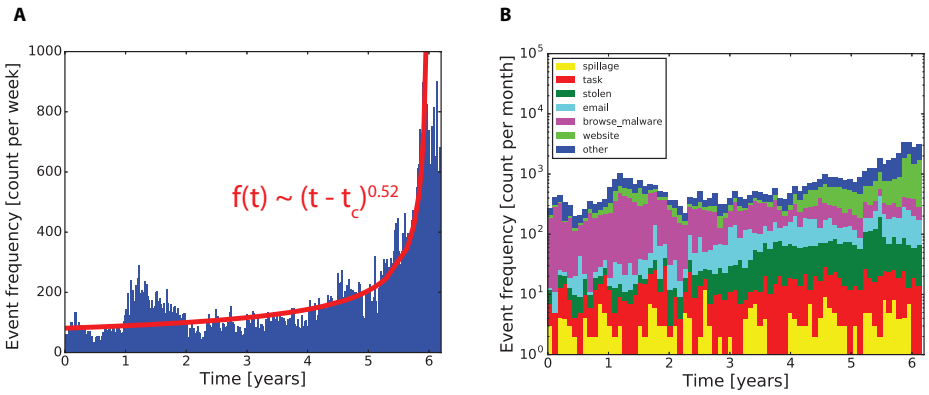


**Fig. 1. A.** Evolution of event count (binned per week). The increase follows a super-exponential growth with finite-time singularity given by equation 2 with $\nu = 0.52$ and $t_c = 6$. **B.** Evolution of event count (binned per month) per category of events (cumulative). While the first half period was dominated by incident associated with malware from Web browsing and USB (purple), the second half period exhibits the rise of stolen devices (dark green) and incidents associated with websites (light green).

Cyber security and cyber threats are rapidly evolving, with new vulnerabilities announced on a daily basis. Over the six years of our data set, changes in security safeguards, network structure, and security processes have occurred. However, the distribution of effort is consistently well accounted for by a power law model (see Table 1). Furthermore, we find evidence that all incidents taken together are overall becoming less extreme over time. Figure 2A shows the distribution of all events greater than 2 man-hours over the time periods from 1 to 12. The exponent $\mu$ governing the skewness of the tail distribution of the power law for the severity of incidents is generally increasing over time (meaning that incidents are becoming less extreme). This suggests that the organization has become more efficient at dealing with large cyber incidents. The CCDF represented in Figure 2A shows that the distribution becomes less heavy-tailed (and thus less extreme) over time. For typical frequencies 10%, 1%, and 0.1% from the main tail distribution (i.e., disregarding outliers), effort has been reduced by respectively, 2, 40 and 600 man-hours in the six-year time frame considered (i.e., between periods 1 and 12) and up to statistical fluctuations as reported in Table 1.

A comparison of small versus large incidents shows that over time, more resources are being devoted to smaller incidents. For instance in period 5, there are more small events (1,026) than tail events (688); yet, the tail events generated over 5 times more investigation hours. In presence of outliers, the aggregated costs are even more skewed towards extreme values. Consider period 8** in which 3,355 events generated 11,774 hours (491 days) of work: the 14 most extreme events (each having required more than 41 hours of work) account for 5,284 hours (i.e., 220 days), and among them, the most

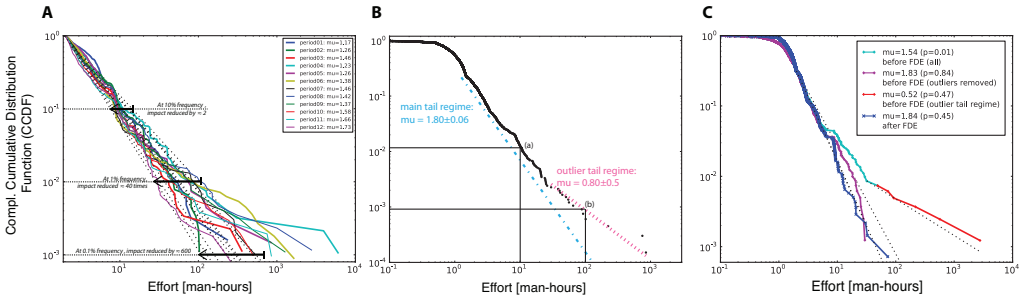


**Fig. 2. A.** Evolution of the complementary cumulative distribution function (CCDF) of effort for all incident types. **B.** The CCDF of investigation time for all incidents in the 11th time period. Note the three regimes: small incidents (here for effort <0.5 man-hours), the main tail (blue dotted line), and an outlier tail (magenta dotted line). **C.** Case study: implementation of Full-Disk Encryption (FDE) and its dramatic effect on the tail distribution.

| type | Period | count events | count events (cost < xmin) | count events (cost >= xmin) | sum events | sum cost (cost < xmin) | sum cost (cost >= xmin) | xmin | largest event | exponent | p-value | std-error | outliers |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| All | 1 | 1'112 | 489 | 623 | 4'588 | 493 | 4'095 | 2 | 233 | 1.17 | **0.91** | 0.08 | - |
| All | 2 | 2'581 | 1'673 | 908 | 7'451 | 1'751 | 5'700 | 2 | 103 | 1.27 | **1.00** | 0.07 | - |
| All | 3 | 3'902 | 3'005 | 897 | 8'391 | 3'123 | 5'268 | 2 | 358 | 1.46 | **0.38** | 0.07 | - |
| All | 4 | 3'457 | 2'535 | 922 | 19'264 | 2'596 | 16'668 | 2 | 6'010 | 1.22 | 0.04 | 0.07 | 2 outliers: 3928 + 6010 = 9938 |
| All | 4* | 3'455 | 2'535 | 920 | 9'326 | 2'596 | 6'729 | 2 | 307 | 1.24 | 0.83 | 0.07 | |
| All | 5 | 1'714 | 1'026 | 688 | 6'257 | 1'025 | 5'231 | 2 | 842 | 1.27 | **0.26** | 0.08 | - |
| All | 6 | 2'716 | 1'592 | 1'124 | 12'130 | 1'591 | 10'540 | 2 | 1'640 | 1.40 | **0.09** | 0.06 | - |
| All | 7 | 2'291 | 1'305 | 986 | 8'126 | 1'433 | 6'694 | 2 | 525 | 1.45 | **0.26** | 0.06 | - |
| All | 8 | 3'355 | 2'501 | 854 | 11'774 | 2'261 | 9'514 | 2 | 2'762 | 1.40 | 0.03 | 0.07 | outliers regime (c.f. next 2 lines) |
| All | 8* | 3'341 | 2'501 | 840 | 6'490 | 2'260 | 4'229 | 2 | 40 | 1.53 | 1.00 | 0.07 | - |
| All | 8** | 3'355 | 3'341 | 14 | 11'774 | 6'490 | 5'284 | 41 | 2'762 | 0.72 | 0.50 | 0.53 | - |
| All | 9 | 3'774 | 2'879 | 895 | 10'317 | 2'840 | 7'478 | 2 | 1'267 | 1.34 | **0.14** | 0.07 | - |
| All | 10 | 5'922 | 4'892 | 1'030 | 11'017 | 4'685 | 6'332 | 2 | 501 | 1.59 | **0.15** | 0.06 | - |
| All | 11 | 6'632 | 5'595 | 1'037 | 12'299 | 5'457 | 6'842 | 2 | 850 | 1.67 | 0.05 | 0.06 | outliers regime (c.f. next 2 lines) |
| All | 11* | 6'618 | 5'595 | 1'023 | 9'799 | 5'457 | 4'342 | 2 | 30 | 1.80 | 1.00 | 0.06 | - |
| All | 11** | 6'632 | 6'618 | 14 | 12'298 | 9'799 | 2'500 | 30 | 850 | 0.89 | 0.53 | 0.53 | - |
| All | 12 | 20'192 | 18'545 | 1'647 | 26'961 | 18'594 | 8'366 | 2 | 313 | 1.72 | **0.25** | 0.05 | - |

**Table 1.** For large events, the power law model cannot be rejected ($p$-value <0.10), with the exception of periods 4, 8, and 11 (highlighted in orange), for which extreme outliers were detected. In period 4, two points deviate from the main tail regime, while for 8 and 9, an outlier tail regime is made up of several points, whose statistical properties are also best explained by a power law. For periods 4, 8 and 11, we provide a corrected model to account for the main tail (*) and outlier tail regime (**).

extreme event alone accounts for 2,762 hours (i.e., 115 days) of work.

Historical cyber incident data can also be used to validate empirically some security investment decisions [6]. The organization studied here implemented full disk encryption (FDE), which increases the difficulty of extracting information from lost or stolen devices by requiring a password before a computer boots. Most data breach notification laws in the US States do not require breach notification when stored information was fully encrypted. We observed and measured the effects of FDE on lost-device investigation times by observing that the outliers disappear after the FDE policy was implemented (see Figure 6). The whole pre-policy distribution is not well fitted by a power law tail (cyan + red lines, $p$ <0.10 implies that we cannot reject the null hypothesis that the sample is not drawn from a power law) because an outlier regime exists (red

line, $\mu = 0.52$, $p = 0.46$). Comparing the distribution of severity before FDE policy implementation without the outliers (purple line, $\mu = 1.89$, $p = 0.79$), and the distribution post-FDE implementation (blue line, $\mu = 1.85$, $p = 0.43$), one cannot reject the hypothesis that they are both drawn from the same distribution. In other words, we have some evidence that the FDE policy has removed the outlier regime, which accounted for 45% of the overall time spent on lost devices. Based on a careful qualitative assessment of descriptions provided for outliers by security engineers at the organization, we found that devices that contain sensitive information (such as personally identifiable information, trade secrets, or other intellectual property) are not very common, but require an overwhelming amount of work, associated with investigation (e.g. analyzing backups, performing forensics, and/or notifying individuals of the breach). Assuming that the outlier tail regime has disappeared, we can state that the time spent on stolen devices has been reduced by a factor 4 (due to the removal of outliers), meaning that full disk encryption in that case has been a very cost-effective prevention measure.

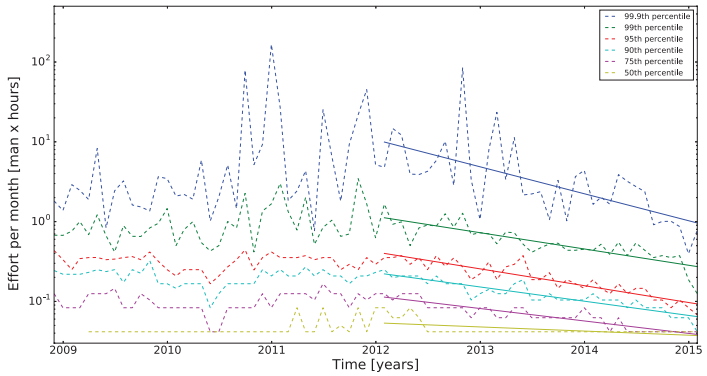


**Fig. 3.** Evolution of the effort required (binned per month) considering percentiles 50, 75, 90, 95, 99, 99.9 (dashed lines) in logarithmic scale on the y-axis. Percentiles provide a measure of the evolution of *extremeness*, which appears here to decrease starting from 2012 for all percentiles above 50. The straight lines exhibit the exponential decays of the higher percentiles as detailed in Table 2.

We finally aim to calibrate the decay of large events over the period 2012 until 2015. Figure 3 shows the evolution of the effort required (binned per month) considering percentiles 99.9, 99, 95, 90, 75, 50 (dashed lines) in logarithmic scale on the y-axis. All high percentiles above percentile 50 decay following exponential functions,

$$N(t) \sim e^{-\Lambda t} \,, \tag{3}$$

with $\Lambda$ the decay rates detailed in Table 2. The excellent fits of exponential decays of high percentile events is striking. They provide a measurable and predictable learning

curve for the organization under scrutiny. This learning curve shall be useful to predict the increase of absorption capacity in the future.

| **Percentile** | **decay rate [1/days]** | $p$ | $R^2$ |
|---|---|---|---|
| 99.9 | 0.34 | 0.000 | 0.43 |
| 99 | 0.20 | 0.000 | 0.71 |
| 95 | 0.21 | 0.000 | 0.83 |
| 90 | 0.18 | 0.000 | 0.83 |
| 75 | 0.16 | 0.000 | 0.83 |
| 50 | 0.05 | 0.000 | 0.33 |

**Table 2.** Exponential decay rates for percentiles above 50.

### 5.2 Recurrence intervals of events

At the aggregate level and over the whole time period, we investigate recurrence intervals, which are a very common measure in quantitative risk management [39], in particular for natural disasters [47] such as earthquakes [48]. Recurrence intervals are the expected time before an event of given size repeats. Here, the rate of all incidents increases dramatically (see Figure 1A). Yet, most of this increase is driven by events that require little effort (i.e., less than 2 hours). For larger events, the recurrence intervals remain overall stable. We find that the distribution of waiting times between two events of effort thresholds $\mathbf{S} = \{6, 12, 24, 48, 168, 720\}$ man-hours follows a exponential distribution,

$$P(T > t) \sim e^{-\lambda t} \tag{4}$$

which describes a memoryless Poisson point process with mean return rate $\beta = 1/\lambda$ [47]. For all effort thresholds $s$, we tested the Poisson process hypothesis. For that, we calibrated the exponential distribution of waiting times between two events using ordinary least square of $log(P)$ as a function of $t$. For all effort thresholds, the Poisson process hypothesis cannot be rejected with high significance (see Table 3).

The results in Table 3 show that while an event of 6 man-hours effort returns on average every 3 days, an event of 24 man-hours returns on average every 24 days. Similarly, an event of seven days ($24 \times 7 = 168$ man-hours) returns on average every 5.5 months, while an event of 100 days returns on average every 16.6 months. These results may contradict the above results, which show an exponential decay of events with percentiles above the median. Here, calculating the return times was made over the whole 2008 until 2015 period, and extreme events represent a minority of events and influence only marginally the overall recurrence of events above much lower thresholds $\mathbf{S}$.

| Effort [man-hours] | Recurrence intervals [days] | $p$ | $R^2$ |
|---|---|---|---|
| $> 6$ | 2.99 | 0.000 | 0.98 |
| $> 12$ | 8.02 | 0.000 | 0.99 |
| $> 24$ | 24.17 | 0.000 | 0.91 |
| $> 48$ | 41.87 | 0.000 | 0.99 |
| $> 168$ | 153.91 | 0.000 | 0.95 |
| $> 720$ | 465.97 | 0.000 | 0.96 |

**Table 3.** Return times for events with effort $> 6$ man-hours. For these thresholds, the memoryless Poisson process hypothesis cannot be rejected with very high confidence (in all cases $p < 0.001$ and $R^2 > 0.90$).

### 5.3 Resolution time, effort & excess effort

Another key aspect of the cyber incident dynamics is the relation between resolution time, effort, and excess effort (effort above the baseline). Figure 4 shows the relation between resolution time $\Delta T$ and effort $S$. Extreme events [percentiles $99.9$ and $99.99$ (resp. yellow and black dashed lines)] are around or above the line $S = \Delta T$, which means that large events tend to be resolved with an effort equivalent to having one person working $24 \times 7$ hours, or having several security officers occupied full time during office hours to resolve the issue. The less effort is required, the even more likely the resolution time will largely exceed required effort. This shows that security officers tend to give highest priority to extreme cyber security incidents. On the contrary, the less effort they require, the more resolution time may be delayed.

We also aim to investigate the dynamics of cyber security event arrival and resolution times. Unfortunately, for 22,620 incidents (out of 60,767 incidents) the resolution time (i.e., when the ticket is closed) is missing. For these events, effort was recorded (in man-hours), suggesting that resolution occurred, even though tickets were not formally closed. These events for which tickets are not formally closed all start around November 2012, when new reporting guidelines were introduced at the organization, mandating opening a ticket for any cyber security incident or event. These guidelines have created extra administrative reporting burden, in particular for events that required little effort. This systematic lack of data accounts for nearly a third of the data set, which is significant if we want to understand workload at the aggregate level. We therefore assume that these incidents were actually resolved, however not formally closed, and we reconstruct the missing *resolution time* by bootstrapping. For that purpose, we employ a non-parametric approach and we randomly sample *resolution time* values from the subset of the data with both *effort* and *resolution time* available. This is equivalent to drawing values from a random variable defined by the probability of resolution time given effort $P(\Delta T|S)$.

Figure 5 exhibits the dynamics of new cyber security events (red), resolved events (green), as well as the weekly mean of resolution times $\Delta T$ (blue). The dynamics of event resolution follow closely the dynamics of new event arrival [Spearman rank cor-

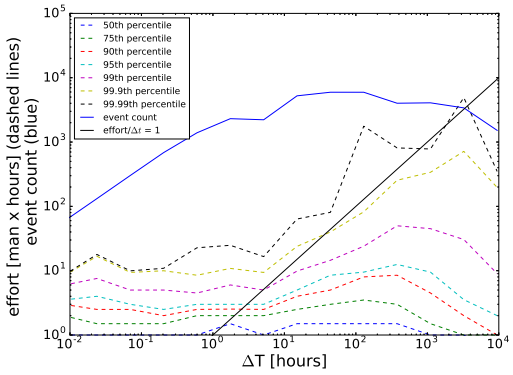


**Fig. 4.** Relation between effort $S$ and resolution time $\Delta T$ in double logarithmic scale from all events for which both $S$ and $\Delta T$ were recorded (i.e., 38,147 events). The black continuous line shows the relation $S = \Delta T$ as a reference. The dashed lines show the percentiles (50th and above) effort as a function of resolution time. The blue line shows event counts as a function of resolution time. Most events have resolution times between 10 and 100 hours.

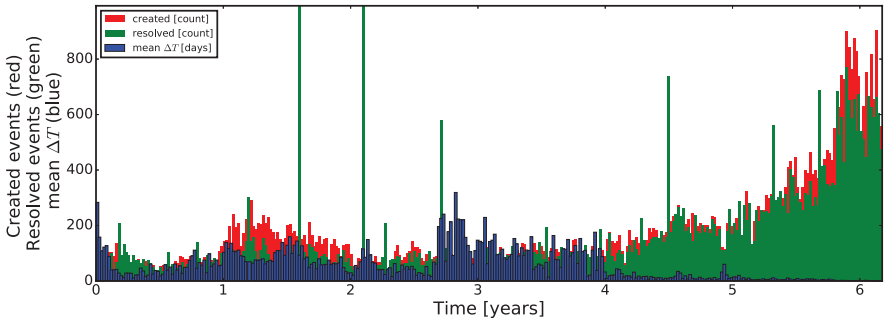


**Fig. 5.** Time series (weekly binned) of new cyber security events (red), resolved events (green), as well as the weekly mean of resolution times (blue). The dynamics of event resolution follow closely the dynamics of new event arrival [Spearman rank correlation $\rho = 0.799$ ($p < 0.001$)]. The resolution times exhibit an increase in the second year and at the end of the third and during the fourth year, and decreases sharply from the fifth year.

relation $\rho = 0.799$ ($p < 0.001$)]. The mean resolution times exhibit an increase in the second year, as well as at the end of the third year. It decreases sharply from the fifth year. This sharp reduction of resolution time (as well as effort per event) is one the central explanation for the unique capacity building, which was required to absorb a super-exponential increase of cyber security events.

We aim to bring fine-grained insights on the hourly dynamics of cyber security events, and how much effort they require at the aggregate level when they are summed together over time. For that, we assume that effort is uniformly distributed over the resolution period. Therefore, the instant effort is given by $S/\Delta T$ over the period from $T_{open,i}$ until $T_{resolved,i}$ for any event $i$. The total instant effort is obtained by summing all uniform efforts over all events. Note that our method to measure instant effort overlooks the typical daily, weekly and yearly seasonalities associated with human activity. We assume 24x7 SOC activity. Figure 6 show instant effort (binned hourly) for January 2012 (Figure 6A) and for year 2012 (Figure 6B). The monthly and yearly scales were chosen to best exhibit the features of instant effort. These panels reflect the whole period under scrutiny. Instant effort exhibits a stable baseline activity, decorated with spiky patterns. To better understand the time series and its regularities, we shall separate both features. To efficiently separate baseline activity from large deviations the smoothing is performed by computing the median activity in the 24-hour centered window. The baseline instant effort is extremely stable over the 6.2 year study period with median effort 1.82 men per hour (5th and 95th percentiles baseline instant effort respectively equals to 0.72 and 3.49 men per hour).

We now turn to excess effort beyond the baseline. In order to absorb sudden changes of effort triggered by the arrival of a new event, it is essential to quantify the *excess effort* beyond the baseline effort. For that, we consider the difference between the instant effort (red) and the smoothed instant effort (blue) as a random variable, from which a value is drawn every hour. Figure 6C shows the distribution of positive variations (red), which follows a power law tail with exponent $\alpha_{pos} = 1.41(4)$ obtained by Maximum Likelihood Estimator (MLE). As a consequence, there is 3.8% (resp. 0.15%) chance that an excess variation of 10 (resp. 100) men per hour will occur within one hour. We further tested dependence between excess effort and baseline effort. We found no statistically significant dependence. As a result, it seems that excess effort is not related to high or low baseline effort. For completeness, negative variations are also shown on Figure 6C. Because the distribution of instant effort is skewed and the smoothing is performed with median values, it is expected that negative variations shall be highly centered, which is indeed what we observe [power law tail with $\alpha_{neg} = 2.98(9)$ obtained by MLE].

### 5.4 Event frequency dependence across sites

The organization under scrutiny has multiple sites (resp. subsidiaries). These sites have various degree of autonomy regarding cyber security incidents. No qualitative information is available for that matter. It is of importance to understand how different sites of a same organization differently or similarly absorb cyber security events. We follow a simple *portfolio management* approach commonly used to measure aggregate risks of portfolio financial products [39]. For that, we measure rank correlations of time series of event counts (weekly bins) between 15 sites with more than 250 events over 6.2 years. Figure 7 shows the correlation matrix across sites, with sites ordered by decreasing number of events (see Table 4). The 9 sites with most cyber security events exhibit high dependence (Spearmann rank correlation $0.24 < \rho < 0.60$), in particular sites $1, 2, 3$ in comparison with site $0$ ($\rho \geq 0.50$). These results suggest that sites share common cyber

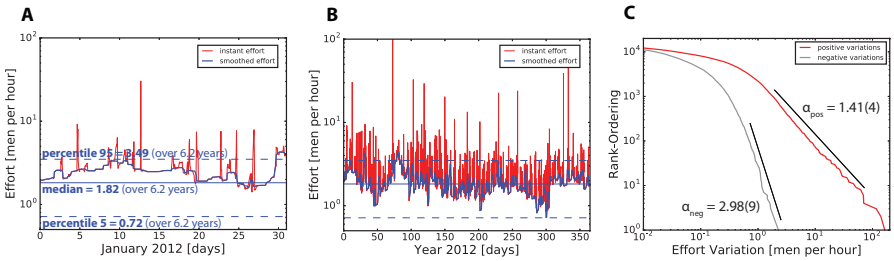


**Fig. 6. A.** Hourly (red) and smoothed effort during January 2012. Smoothing is performed by taking the median activity on a rolling window of 24 hours. The median ensures that baseline activity is well separated from the large *spiky* excursions of positive effort variations. The smoothed effort exhibits high stability over the 6.2 years of data with media, 5th and 95th percentile effort respectively equal to 1.82, 0.72, 3.49 men per hour **B.** Similar time series for the whole year 2012, which is comparable to the whole period of the data set. **C.** Rank-ordering (i.e., unnormalized complementary cumulative distribution function) of effort variation above (resp. below) the smoothed effort. Both distributions exhibit power law tail distributions with exponents resp. $\alpha_{pos} = 1.41(4)$ ($x_{min} = 1$ and $p = 0.82$) and $\alpha_{neg} = 2.98(9)$ ($x_{min} = 0.4$ and $p = 0.82$). Both fits were obtained with Maximum Likelihood Estimator (MLE) and confidence intervals were obtained by bootstrapping following [8, 10, 40].

security incidents, which may hit all sites at the same time. High dependence may also be associated with top-down propagation of cyber security notifications, or rather with horizontal information sharing and coordination across sites.

| **site** | **0** | **1** | **2** | **3** | **4** | **5** | **6** | **7** | **8** | **9** | **10** | **11** | **12** | **13** | **14** |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **events** | 20,459 | 7,252 | 5,026 | 4,997 | 4,659 | 4,515 | 4,233 | 2,577 | 1,871 | 1,100 | 1,045 | 867 | 782 | 751 | 386 |

**Table 4.** Number of events recorded per site over the entire period of 6.2 years.

## 6 Discussion

We have analyzed 60,767 cyber security events and incidents recorded at a large organization over a six-year and two months period. In addition, we had access to the effort required to overcome each event (in man-hours).

Our study has unveiled a number of stylized facts associated with the flow cyber security incidents and their resolution by the security operations center (SOC). We first found a super-exponential increase of events recorded in the tracking system. Some categories of events were relatively more prevalent in the first years, and then replaced by other types of security incidents. While the number of events recorded increased,

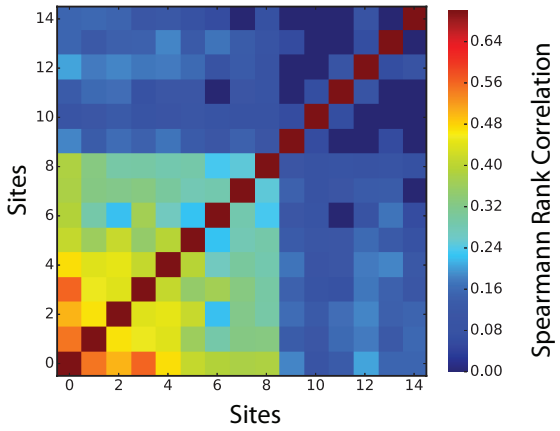


**Fig. 7.** Spearman rank correlation weekly event count between the 15 sites with more than 250 recorded cyber security events over 6.2 years.

we found that human resources devoted to cybersecurity incidents has not significantly increased over time on average (1.82 men per hour on a 24x7 basis). In a nutshell, it appears that the organization under scrutiny has managed to build capacity super-exponentially at nearly constant human resources at the SOC.

We investigated the origins of this dramatic performance increase and we found that the organization managed to reduce the frequency of large and extreme *tail* events. We brought a variety of perspectives to test this hypothesis. We first noticed that the overall tail distribution of cyber security incident resolution effort is overall best modeled with a power law distribution with exponent $\mu \approx 1.5$, yet increasing over time (i.e., the distribution is getting less extreme). The tail distribution was periodically decorated with so-called *dragon-king* large outliers [38, 49]. Each of these outlier events may amount to large proportions of effort over the same period. These outliers were successfully tackled by the organization. For instance, we could show how the implementation of full-disk encryption durably removed risks associated with data spillage. We complemented the study of the evolution of the tail distribution, by a non-parametric longitudinal study of percentiles above the median (i.e., large and extreme events), and we found that from 2012 on, they decreased exponentially. The decrease of extreme event frequency may have created capacity to manage more events of smaller size, following a positive feedback loop, which may explain the super-exponential increase of events (see Figure 1A).

We have then investigated the fine-grained dynamics associated with the hourly management of cyber security events. We found that the baseline of human resource effort remained very stable between 0.72 men per hour (5th percentile) and 3.49 men per hour (95th percentile) over the 6.2 years. Yet, the hourly effort is decorated by jumps which size is a random variable. The random variable is well represented by a power

law tail with exponent $\alpha_{pos} = 1.41(4)$. The statistical properties of the tail distribution of jump size imply that the mean converges while the second moment diverges as $n \to \infty$ [47]. These *excess effort* jumps require thorough quantification, and shall be taken into account for cyber security incident response. Absorption of larger shocks requires preparedness and adequate provision of human resources for emergencies in addition to baseline effort. The human resource effort may suddenly jump to 10 men per hour (with probability 3.8% per hour) or to 100 men per hour (with probability 0.1% per hour).

We have finally touched upon portfolios of cyber security incidents. As the organization under scrutiny has several sites (resp. subsidiaries), we could compute the dependence of event frequencies between sites with a sufficient number of events (i.e. $\geq 250$). We have no knowledge whether the organization under scrutiny has a centralized versus a decentralized management of cyber security incidents. We could find that the frequency of events on the site with most events influences more than 50% of the frequency of events in the 3 next sites by ranking order in terms of event frequency. Also, the 9 sites with most events are overall highly dependent, suggesting some form of central organization or at least a common information sharing channels. The 6 remaining sites are much less dependent between each other and with the 9 sites receiving most events. This suggest that these sites are more autonomous regarding the management of their cybersecurity. They may also be less exposed. Or on the contrary, the may be managed with less efficiency. Much more investigation is certainly needed to build a solid risk portfolio theory for cyber security, but we can already draw some relevant conclusions from this prototype regarding organization resilience.

Risk accumulation is one of the top topics in the nascent, yet fast-growing, cyber insurance industry, and organizations shall be concerned similarly at their own scales. A large organization typically holds tens or even hundreds of subsidiaries. Let's consider that cyber security is partly centrally managed and partly locally managed. The level of autonomy of subsidiaries may depend on the efficiency adjusted to the risk portfolio. Both efficiency and risk portfolios can be adjusted in several ways, but there is usually a trade-off. On the one hand, by delegating security, one may get more locally optimized response, but may risk inconsistent cyber threat response with potential implications depending on information system interconnectedness between subsidiaries. On the other hand, centrally managed cyber security may increase dependence between subsidiaries, and generate long-term fragility and difficulties in provisioning response resources. Similarly, more centrally managed information systems bring more operational efficiency. However, they are certainly more homogeneous and thus, they may generate more dependence of cyber security events across subsidiaries.

The risk portfolio management approach also demonstrates that measuring the dynamics of effort to overcome cyber security incidents encapsulates information, not only about the technical challenges associated with cyber security, but also on how humans in charge can handle them. Behind the time and effort spent on solving a security event, there are humans with various levels of technical background, expertise, eager-

ness to learn with agility and ability to recognize outstanding challenges. Again, more research is needed to bring further supporting evidence, but it is striking to see how the security officers at the organization studied here seem to have produced impressive return-on-scale. We believe that learning is one aspect, but quite surely capacity to use programmatic approaches (e.g., full-disk encryption) to durably reduce effort associated with tackling events of similar nature, certainly has played a key role in building capacity. We shall therefore propose that security officers as individuals, but also organized teams, as well as their capacity to deploy programmatic responses to security challenges should receive much more research attention. In other words, the human factor is not only related to users, but primarily to the cognitive capabilities, expertise, incentives and psychology of the very humans in charge of cyber security.

Finally, we would to stress that aggregates through statistics have immediate advantages for managing cyber risks. For instance, the risk portfolio approach allowed comparing different sites without knowing any detail on the technical aspects of cyber security events. Instead of subsidiaries, the same principle could be used to manage cyber risks of a portfolio of suppliers at the aggregate level, without knowing in much details the technical operations associated with cyber security. Statistics may be privacy preserving in this case. Many cybersecurity practitioners recognize the lack of common language and dialogue between executives, in particular the Chief Risk Officer, and technical cyber security teams. We propose that knowing how security officers handle cyber security events provides the essential part of information needed to take appropriate management and risk provisioning decisions, while avoiding micro-management. Measuring performance from ticketing systems, may however require to set proper incentives and clear guidelines to avoid reporting moral hazard and other biases.

### 6.1 Limitations and future work

In this exploratory work, we believe we have only scratched the surface of the potential associated with studying the dynamics of human resource effort devoted to cyber security events. The three main limitations are (i) theoretical, (ii) empirical, which in turn does not allow to deepen on (iii) organization design challenges. From a theoretical perspective, we first still lack a general mechanism, which would explain the reduction of large and extreme events and how shock absorption increases. Second, a proper theory of cyber security incident portfolio management shall be developed. Once overcome, these two theoretical challenges deserve further empirical testing, possibly with additional and more complete data from other organizations. For instance, our data only the human resource effort for incident (resp. event) resolution, but does not consider potential additional costs, such as operation disruption or litigation. We also could not get a clear sense of urgency associated with events. In the future, we could for instance envision running sentiment analysis from natural language processing of event descriptions.

Establishing a benchmark between organizations is also one of our future work priorities. Is the organization studied here of its own kind or our findings generalize? What difference shall we expect within and across industries, or perhaps regarding physical

proximity, software proximity (use of similar software) or homogeneity of talents recruited across organizations?

We also believe that quantitative results that depict the individual and collective behaviors of security officers at such fine-grained level has the potential to inspire further targeted research regarding incentives, cognition, and psychology on this group of individuals.

## 7 Conclusion

Organizations are constantly subjected to cyber attacks, incidents, and other cyber security events. Absorbing the flow of incidents at sustainable costs is paramount to maximize resilience. Here, we have uncovered a surprising virtuous circle, as the organization under scrutiny has successfully managed to reduce the frequency of large and extreme efforts that require the mobilization of important human resources. These significant gains have allowed managing a super-exponential increased number of incidents and events. At a finer-grained level, we could quantify the stable baseline human resource effort deployed by the organization, as well as the probability of *excess instant effort* as a power law tail distribution with exponent close to 1.4. These results provide precisely quantified statistical measures, which could be used to optimize organization design regarding cyber risks. Considering that the organization under scrutiny had several sites (resp. subsidiaries) to which events could be attributed, we have developed a prototype of cyber risk portfolio and we have discussed how this approach may significantly improve cyber risk management, namely by using quantitative risk portfolio models. If sufficiently easy to deploy, these quantitative models may replace qualitative approaches in the future. We believe that the approach presented here is simple enough, yet highly informative, and thus, may be broadly adopted in a not so distant future.

## References


1. Florêncio, D., Herley, C.: Sex Lies and Cyber-Crime Surveys. In: Economics of Information Security and Privacy III. Springer Science + Business Media (jul 2012) 35–53
2. Tøndel, I.A., Line, M.B., Jaatun, M.G.: Information security incident management: Current practice as reported in the literature. Computers & Security **45** (sep 2014) 42–57
3. Arora, A., Hall, D., Piato, C., Ramsey, D., Telang, R.: Measuring the risk-based value of IT security solutions. IT Professional **6**(6) (nov 2004) 35–42
4. Pisarenko, V.: Non-linear growth of cumulative flood losses with time. Hydrological Processes **12**(3) (1998) 461–470
5. Sachs, M., Yoder, M., Turcotte, D., Rundle, J., Malamud, B.: Black swans, power laws, and dragon-kings: Earthquakes, volcanic eruptions, landslides, wildfires, floods, and soc models. The European Physical Journal Special Topics **205**(1) (2012) 167–182
6. Sornette, D., Maillart, T., Kröger, W.: Exploring the limits of safety analysis in complex technological systems. International Journal of Disaster Risk Reduction **6** (2013) 59–66
7. Chernov, D., Sornette, D.: Man-made catastrophes and risk information concealment. Springer (2016)

8. Maillart, T., Sornette, D.: Heavy-tailed distribution of cyber-risks. The European Physical Journal B **75**(3) (apr 2010) 357–364
9. Edwards, B., Hofmeyr, S., Forrest, S.: Hype and heavy tails: A closer look at data breaches. Journal of Cybersecurity **2**(1) (2016) 3–14
10. Wheatley, S., Maillart, T., Sornette, D.: The extreme risk of personal data breaches and the erosion of privacy. The European Physical Journal B **89**(1) (jan 2016)
11. Cezar, A., Cavusoglu, H., Raghunathan, S.: Outsourcing Information Security: Contracting Issues and Security Implications. Management Science **60**(3) (mar 2014) 638–657
12. Bolot, J., Lelarge, M.: Cyber Insurance as an Incentive for Internet Security. In: Managing Information Risk and the Economics of Security. Springer Science + Business Media (dec 2008) 269–290
13. Ralston, P., Graham, J., Hieb, J.: Cyber security risk assessment for SCADA and DCS networks. ISA Transactions **46**(4) (oct 2007) 583–594
14. Anderson, R., Moore, T.: The Economics of Information Security. Science **314**(5799) (oct 2006) 610–613
15. Frei, S., Schatzmann, D., Plattner, B., Trammell, B.: Modeling the Security Ecosystem - The Dynamics of (In)Security. In: Economics of Information Security and Privacy. Springer Science + Business Media (2010) 79–106
16. Frei, S., May, M., Fiedler, U., Plattner, B.: Large-scale vulnerability analysis. In: Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense - LSAD '06, Association for Computing Machinery (ACM) (2006)
17. Zhang, S., Caragea, D., Ou, X.: An Empirical Study on Using the National Vulnerability Database to Predict Software Vulnerabilities. In: Lecture Notes in Computer Science. Springer Science + Business Media (2011) 217–231
18. Liu, D., Wang, X., Camp, J.: Game-theoretic modeling and analysis of insider threats. International Journal of Critical Infrastructure Protection **1** (dec 2008) 75–80
19. Liu, D., Wang, X., Camp, L.J.: Mitigating Inadvertent Insider Threats with Incentives. In: Financial Cryptography and Data Security. Springer Science + Business Media (2009) 1–16
20. D'Ambros, M., Lanza, M., Robbes, R.: An extensive comparison of bug prediction approaches. In: 2010 7th IEEE Working Conference on Mining Software Repositories (MSR 2010), Institute of Electrical and Electronics Engineers (IEEE) (may 2010)
21. Kjaerland, M.: A taxonomy and comparison of computer security incidents from the commercial and government sectors. Computers & Security **25**(7) (oct 2006) 522–538
22. Arora, A., Telang, R., Xu, H.: Optimal policy for software vulnerability disclosure. Management Science **54**(4) (2008) 642–656
23. Arora, A., Telang, R.: Economics of software vulnerability disclosure. IEEE Security and Privacy Magazine **3**(1) (jan 2005) 20–25
24. Neuhaus, S., Zimmermann, T., Holler, C., Zeller, A.: Predicting vulnerable software components. In: Proceedings of the 14th ACM conference on Computer and communications security - CCS '07, Association for Computing Machinery (ACM) (2007)
25. Shin, Y., Meneely, A., Williams, L., Osborne, J.A.: Evaluating Complexity Code Churn, and Developer Activity Metrics as Indicators of Software Vulnerabilities. IEEE Transactions on Software Engineering **37**(6) (nov 2011) 772–787
26. Rhodes, C., Anderson, R.: Epidemic Thresholds and Vaccination in a Lattice Model of Disease Spread. Theoretical Population Biology **52**(2) (oct 1997) 101–118
27. Chew, E., Swanson, M., Stine, K.M., Bartol, N., Brown, A., Robinson, W.: Performance measurement guide for information security. Technical report (2008)
28. Nozaki, H., Tipton, M.: Information Security Management Handbook Sixth Edition, Volume 5. Informa UK Limited (sep 2011)

29. Ahmad, A., Maynard, S.B., Shanks, G.: A case analysis of information systems and security incident responses. International Journal of Information Management **35**(6) (dec 2015) 717–723
30. Condon, E., He, A., Cukier, M.: Analysis of Computer Security Incident Data Using Time Series Models. In: 2008 19th International Symposium on Software Reliability Engineering (ISSRE), Institute of Electrical and Electronics Engineers (IEEE) (nov 2008)
31. Joh, H., Malaiya, Y.K.: Seasonal Variation in the Vulnerability Discovery Process. In: 2009 International Conference on Software Testing Verification and Validation, Institute of Electrical and Electronics Engineers (IEEE) (apr 2009)
32. Bando, K., Tanaka, K.: Trend Analyses of Accidents and Dependability Improvement in Financial Information Systems. In: 2011 IEEE 17th Pacific Rim International Symposium on Dependable Computing, Institute of Electrical & Electronics Engineers (IEEE) (dec 2011)
33. Paté-Cornell, M., Kuypers, M., Smith, M., Keller, P., et al.: Cyber risk management for critical infrastructure: a risk analysis model and three case studies. Risk Analysis **38**(2) (2018) 226–241
34. Böhme, R.: Security Metrics and Security Investment Models. In: Advances in Information and Computer Security. Springer Science + Business Media (2010) 10–24
35. Bojanc, R., Jerman-Blažič, B.: Towards a standard approach for quantifying an ICT security investment. Computer Standards & Interfaces **30**(4) (may 2008) 216–222
36. Thomas, R.C., Antkiewicz, M., Florer, P., Widup, S., Woodyard, M.: How bad is it?–a branching activity model to estimate the impact of information security breaches. A Branching Activity Model to Estimate the Impact of Information Security Breaches (March 11, 2013) (2013)
37. McNeil, A.J., Frey, R., Embrechts, P.: Quantitative risk management: Concepts, techniques and tools. Princeton university press (2015)
38. Wheatley, S., Sornette, D.: Multiple Outlier Detection in Samples with Exponential & Pareto Tails: Redeeming the Inward Approach & Detecting Dragon Kings. Swiss Finance Institute Research Paper No. 15-28 (aug 2015)
39. Malevergne, Y., Sornette, D.: Extreme financial risks: From dependence to risk management. Springer Science & Business Media (2006)
40. Clauset, A., Shalizi, C.R., Newman, M.E.J.: Power-Law Distributions in Empirical Data. SIAM Review **51**(4) (nov 2009) 661–703
41. Embrechts, P., Resnick, S.I., Samorodnitsky, G.: Extreme Value Theory as a Risk Management Tool. North American Actuarial Journal **3**(2) (apr 1999) 30–41
42. Embrechts, P., Klöppelberg, C., Mikosch, T.: Modelling Extremal Events. Springer Science + Business Media (1997)
43. Stephens, M.A.: EDF Statistics for Goodness of Fit and Some Comparisons. Journal of the American Statistical Association **69**(347) (sep 1974) 730–737
44. Moffatt, H.: Euler's disk and its finite-time singularity. Nature **404**(6780) (2000) 833
45. Johansen, A., Sornette, D.: Finite-time singularity in the dynamics of the world population, economic and financial indices. Physica A: Statistical Mechanics and its Applications **294**(3-4) (2001) 465–502
46. Easwar, K., Rouyer, F., Menon, N.: Speeding to a stop: the finite-time singularity of a spinning disk. Physical Review E **66**(4) (2002) 045102
47. Sornette, D.: Critical phenomena in natural sciences: chaos, fractals, self-organization and disorder: concepts and tools. Springer Science & Business Media (2006)
48. Saichev, A., Sornette, D.: Theory of earthquake recurrence times. Journal of Geophysical Research: Solid Earth **112**(B4) (2007)
49. Sornette, D., Ouillon, G.: Dragon-kings: mechanisms, statistical methods and empirical evidence. The European Physical Journal Special Topics **205**(1) (2012) 1–26