An Empirical Analysis of Cyber Security Incidents at a Large Organization

Marshall A. Kuypers[1], Thomas Maillart[2], and Elisabeth Paté-Cornell[1]

[1] Department of Management Science and Engineering, Stanford University, Stanford, CA
{mkuypers, mep}@stanford.edu
[2] School of Information, UC Berkeley, Berkeley, California
maillart@berkeley.edu

Abstract. Every day, security engineers cope with a flow of cyber security incidents. While most incidents trigger routine reactions, others require orders of magnitude more effort to investigate and resolve. How security operation teams in organizations should tune their response to tame extreme events remains unclear. Analyzing the statistical properties of sixty thousand security events collected over six years at a large organization, we find that the distribution of costs induced by security incidents is in general highly skewed, following a power law tail distribution. 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. We illustrate this result with a case study focused on the empirical effects of full disk encryption on the severity of incidents involving lost or stolen devices.

Keywords: Cyber security, risk analysis, incident management systems

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. Information sharing initiatives are making progress,[1] 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 these databases are heavily biased towards incidents that require public disclosure and large-scale incidents at major organizations that are reported by the media [3]. Therefore, many open questions exist about the statistical properties of cyber security incidents and their implications for the organization of security response.

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

Fortunately, cyber security incident data are collected for the purpose of auditing, compliance, and management [4]. These data can be leveraged to derive risk metrics, which in turn may lead to more informed security investments [5,6]. Here, we present a statistical analysis of 60,000 cyber security incidents spanning six years that occurred at a large US organization.

Like many other natural and manmade disasters, cyber incidents involve extreme events. Even if these events do not necessarily lead to catastrophic damage, they may have unintended consequences. For instance, at a large Internet company, it was reported to the authors that the Heartbleed vulnerability[2] required a team of more than five expert security engineers over a whole week to safely deploy the necessary patches across the infrastructure, with obvious disruptive effects and delays on the normal workflow of cyber security incident management.[3] If even the best equipped and security aware organizations are struggling with such extreme events, which are relatively common [7], one may question the ability of average organizations facing similar challenges to quickly respond to critical incidents with enough manpower. For all organizations concerned with cyber security incidents, deciding how much in-house human expertise versus outsourcing [8], risk acceptance [5] and transfer (e.g., insurance) [9], is becoming an increasingly pressing question.

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, a single rare event accounts for nearly 30% of the total hours allocated by the security operation center over a period of six months (see e.g., Table 2, period 8). Calibrating the statistical properties of cyber security incidents requires extreme risk modeling, which has been used to address other outstanding cyber security challenges, such as heavy-tailed distributions of personal data breaches [7,10,11]. We find that cyber security incidents tend to become overall less extreme over time, although some large outliers (so-called "Dragon-Kings" [12]) are found to have unique properties compared to less extreme events. We illustrate this latter point with lost or stolen devices (e.g., laptops), and we show how the consequences may be dramatically different depending on the type of information carried on the device. Furthermore, we show how the organization that we studied addressed this challenge by implementing a full disk encryption (FDE) policy.

The rest of the paper is organized as follows. First, we review background research in Section 2. We then present the nature of cyber security incidents contained in our data set in Section 3, followed by a description of the standard statistical methods

[2] See heartbleed.com for more information.

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

employed for the study of extreme risks (Section 4). In Section 5, we report the evolution of incident frequency and severity distribution, and in Section 6 we discuss the results. Finally, we highlight limitations and future work (Section 7) and we present our conclusions (Section 8).

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 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 [17]), personal and sensitive data thefts [7,10,11] as a result of Internet attacks [18], insider threats [19,20], and human mistakes [21]. 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 [22]. In many cases, public release has brought better understanding about the risks associated with these events [23,24], such as robust statistical models of personal data breaches [7,10,11], and predictive algorithms of software vulnerabilities [25-28].

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 a comparable way to vaccination in epidemiology [29]. For the time being, most organizations have no requirement and usually no incentive to disclose information. Nevertheless, recording cyber security incidents at an organization has long been recognized as a critical part of IT security and much work has been written about the appropriate collection of cyber incident data [30-33]. US-CERT requires that certain information be recorded and reported when an incident occurs on a federal information system [22].[4] Therefore, many organizations record data in incident management systems but may not fully leverage that information. Ahmad et al. analyzed incident management systems at a financial institution and found that miscommunication and organizational barriers prevented incident data from being

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

best used [34]. Unfortunately, but for understandable legal and strategic reasons, security incident data are rarely shared. Therefore, academic research is typically limited to either theoretical considerations or surveys about incidents and practice, instead of data-driven empirical analyses [4].

Nevertheless, some investigations have brought insights associated to cyber security incidents within organizations. In 2008, Condon et al. published an analysis of security incidents at the University of Maryland [35]. However, the data consists primarily of malware incidents, and only their frequency was studied but not their severity. In another analysis using University of Maryland data from incidents and intrusion prevention devices, researchers found that the number of attackers or attack signatures did not increase the number of security incidents, but that the number of corrupted computers on the university network did [36]. Others have studied vulnerability disclosures [37], or even the failure of financial information systems [38], but work analyzing cyber incident data currently remains 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 and remediation. Kuypers and Paté-Cornell have used these probabilistic outputs as inputs to quantitative risk models to assess cyber risks in dollar terms by modeling the cost of incident investigation, and also reputation damage, business interruption, intellectual property loss, and other costs [6]. The generation of these probabilistic inputs is critical, given the heavy-tailed nature of some cyber incidents. Other cyber risk models have historically used expected values of losses instead of probabilistic inputs, probably because of data constraints [39, 40]. Models that use Monte Carlo simulations or other methods to probabilistically assess risk provide much more information about cyber risk [41].

3 Data

The data set used for this study is organized by security incident type. Security engineers create tickets and record information including a tracking number, the date of the incident, the number of systems impacted, the total number of hours of investigation, the suspected attackers, and other details about the incident.[5] These incident tracking systems are common in large organizations and are used for workflow

[5] Due to security concerns, the data will not be attributed to the organization. Further, identifying characteristics about the organization have been obfuscated. It is the authors' hope that this paper will serve as an example for other organizations to publish cyber security data for researchers.

tracking, auditing purposes, and have great leverage potential for cyber security intelligence.[6] Note that these are not automatically generated log data, but manually entered tickets.

The data, which span from roughly the beginning of 2009 to the end of 2014, 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.[7] The perpetrators also represent nearly all categories of attackers, including insiders, hacktivists, criminals, and nation states.

Each incident in this data set is characterized by a number that specifies how many man-hours were required to investigate the incident, which generally also 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. For the purposes of this paper, the incidents are classified into the categories listed below. We study each of these categories separately.

a. **Data spillage[8]:** Incidents that possibly disclose information to unauthorized individuals are categorized as data spillage. For example, an employee could forget to encrypt an email that contains social security numbers.

b. **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 ("phishing"), or attach malicious files to an email to infect a user's machine with malware.

c. **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 the type of device and the encryption level.

d. **Tasks:** These involve 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).

---

6 Typically, organizations with a mature cyber security program or security operations center will record incidents. The authors have confirmed the existence of varying degrees of cyber incident tracking systems with dozens of Fortune 500 companies, as well as many government agencies.

7 The incidents are more appropriately described as 'incidents that the security operations center (SOC) deals with' instead of 'cyber attacks against the organization'.

8 Data spillage includes incidents that are often categorized as 'insider attacks'. We do not use the term 'insider attacks' because the majority of incidents are accidental and not malicious.

e. **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.

f. **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.[9]

g. **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.

The data offer a comprehensive view of the human resources deployed by a large organization over a long time period, and 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 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 good measure of the effort required to overcome these incidents.

4 Method

Today, most organizations face a flow of cyber security incidents that occur with some frequency and severity. By nature, each event requires a unique response effort. If the random variable that represents the severity of an event is defined by a narrow and well-centered distribution, the overall cost stems from the frequency of these events, since all incidents have roughly the same severity. In this study, however, we observe that the costs of some events are larger by orders of magnitude than the median incident severity. The severity and the frequency of large events relative to average events is determinant and must be appropriately quantified.

[9] Note that malware delivered via email (e.g. a malicious attachment or an email that contains a link to a malicious website) is classified under email attacks.

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, which require less than 2 hours of work, (ii) heavy-tail large events, and (iii) extreme outliers, which presumably stem from qualitatively different incidents (see Section 5.4 for further details). 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, but also by threat category, as described in Section 3.

4.1 Evolution of Incident Frequency

The frequency is measured by the number of incidents per month, at the aggregate level, above a severity threshold (respectively 3, 10, 20 and 50 man-hours), and by category. At this point, we only consider long-term trends over the full six year period, and find that the rate of events evolves slowly overall. Therefore, we resort to linear regressions to quantify the rate of events.

4.2 Nature and Evolution of Incident Severity

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 [42]. 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 [7] 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.[10] A variety of tools, such as Extreme Value Theory (EVT), have been developed to assess the probability of an extreme event beyond observations [43]. 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 [44].

[10] It shall not be a surprise that first evidence for extreme risks was found for earthquakes, for which historical and geological footprints can be traced over very long time periods.

Here, we observe that event severity follows a power law tail distribution given by,[11]

$$P(S \geq s) = \frac{1}{s^{\mu}}$$

where s is the severity and μ its exponent, within the boundaries defined by the lower threshold, and for values smaller than the outliers. We take the power law model as our default model, and we want to assert within which lower (set at 2 man-hours for this study) and upper boundaries, the null hypothesis (that our data are not drawn from this model) cannot be rejected.

Given a sample, we find the best fit using maximum likelihood estimator (MLE) [10,42]. 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 [45]. 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 model's description of the data. We chose a relatively conservative level ($p > 0.1$) as the rejection level for the null hypothesis [42]. 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 [7]. While more sophisticated methods have been developed to account for extreme outlier regimes [46], these model specifications, encompassing the main and outlier tail regimes, are sufficient to account for the statistical properties found in our data set.

5 Results

5.1 Cyber security incident frequency

The frequency of all cyber incidents is found to be increasing over time (see Figure 1), but the change is driven by small incidents (less than 2 man-hours). The number of large incidents is remarkably constant over time. The rate of incidents taking five hours or more to investigate remains very stable, with an average of 45 incidents per month. The increase in small-scale incidents is attributed to evolving incident recording guidelines over time, which is explained in more detail in Section 5.3.

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

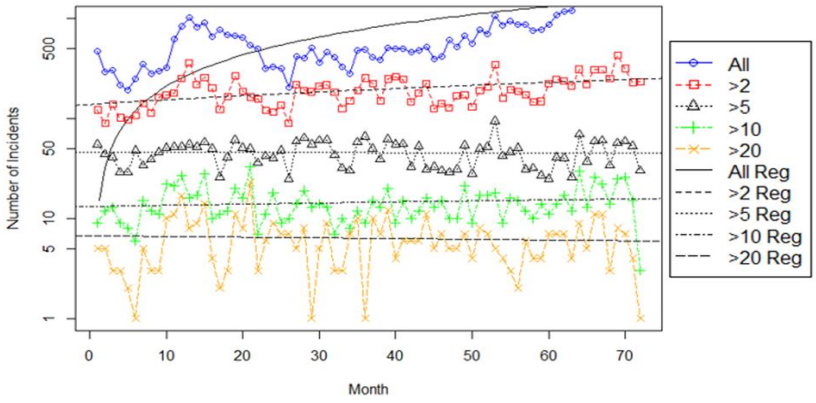

Fig. 1. The number of all cyber security incidents over time, along with linear regressions. Note that the y-axis is presented on a logarithmic scale.

Analyzing the rate of incidents for different attack vectors gives a more detailed view of cyber security at the organization. Figure 2a shows the number of web browsing/ USB incidents recorded over time. The total number of incidents is slowly decreasing over time (2 events per month), while the number of incidents that require more than 10 hours of investigation remains relatively constant over time, at a level of about 7 events per month.

The frequency of lost devices is significantly different from that of browsing malware infections (Figure 2b) as malware incidents occur much more frequently than the loss of devices (thousands of malware incidents per month compared to less than one hundred lost devices per month). Also, the total number of lost device incidents increased dramatically around 2012, from an average of 7 events per month in the first year to 50 events per month in the last year. The increase in the frequency of lost devices is largely due to a change in incident reporting guidelines. Before 2012, losses of mobile phones were often not recorded because those devices did not contain any sensitive information, and losses of security tokens were not recorded because few employees had them.[12] However, as email information began to reside on cellphones, recording stolen or lost devices became more common. Similarly, the number of lost tokens increased as more employees used them. Fitting the last two years of data, during which the reporting guidelines remained consistent, we find that on average, 50 (standard deviation 5) devices were lost per month.

[12] Security tokens are used as a form of two-factor authentication.

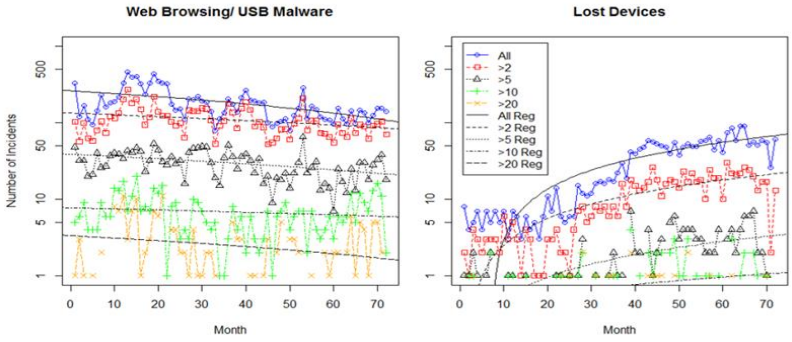

Fig. 2. Web browsing/USB and lost devices incidents over time. Note that the y-axis is on a log scale.

Table 1 shows the rate of incidents for all cyber incident categories, and a linear regression. We find that for this organization, most cyber incidents occur at a surprisingly constant rate.

Table 1. The average monthly number of incidents for each category, and the change per month, calculated using a linear regression. The four vectors with the largest change are highlighted for each division of investigation hours.

	Average (Standard Deviation)					Regression (Standard Error)				
Vector	All Incidents	2+	5+	10+	20+	All Incident	2+	5+	10+	20+
All	797.97 (707.28)	194.33 (67.69)	45.58 (13.32)	14.53 (5.93)	6.33 (3.79)	22.19 (3.0)	1.53 (0.34)	-0.013 (0.076)	0.036 (0.034)	-0.010 (0.022)
Browse	183.90 (90.54)	109.06 (45.40)	29.85 (11.99)	6.83 (4.10)	2.49 (2.67)	-2.13 (0.45)	-0.68 (0.25)	-0.24 (0.062)	-0.025 (0.023)	-0.023 (0.015)
Email	61.07 (51.19)	6.83 (6.46)	2.18 (2.28)	1.03 (1.43)	0.53 (1.14)	1.87 (0.19)	0.13 (0.034)	0.023 (0.013)	0.002 (0.008)	-0.002 (0.007)
Spillage	3.33 (2.28)	2.29 (2.07)	0.96 (1.23)	0.59 (0.82)	0.28 (0.59)	0.022 (0.013)	0.024 (0.011)	0.004 (0.007)	0.001 (0.005)	0.002 (0.003)
Stolen	30.76 (24.97)	10.64 (8.03)	1.72 (1.87)	0.74 (0.92)	0.28 (0.54)	1.05 (0.068)	0.31 (0.027)	0.046 (0.009)	0.010 (0.005)	0.002 (0.003)
Task	10.92 (5.53)	4.014 (2.56)	1.61 (1.41)	0.92 (1.07)	0.57 (0.87)	0.13 (0.028)	0.039 (0.014)	0.013 (0.008)	0.009 (0.006)	0.008 (0.005)
Websites	219.78 (344.08)	19.85 (23.09)	3.72 (2.57)	2.38 (1.69)	1.61 (1.33)	10.43 (1.52)	0.68 (0.10)	0.040 (0.014)	0.012 (0.010)	0.003 (0.008)
Other	225.35 (240.06)	41.65 (30.68)	5.54 (3.46)	2.056 (1.56)	0.58 (0.78)	6.94 (1.1)	1.2 (0.13)	0.1 (0.016)	0.029 (0.008)	0.000 (0.004)

5.2 **Cyber security incident severity**

The impact distribution of cyber incidents is made up of three regimes; low impact incidents (less than 2 man-hours of work), a main tail section, well represented by a

power law distribution, and in some cases, a transient outlier tail regime also well described by a more extreme power law distribution with a smaller exponent. Figure 3 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% (respectively, 0.1%) chance that an event involving more than 10 (respectively 100) hours of work will occur.

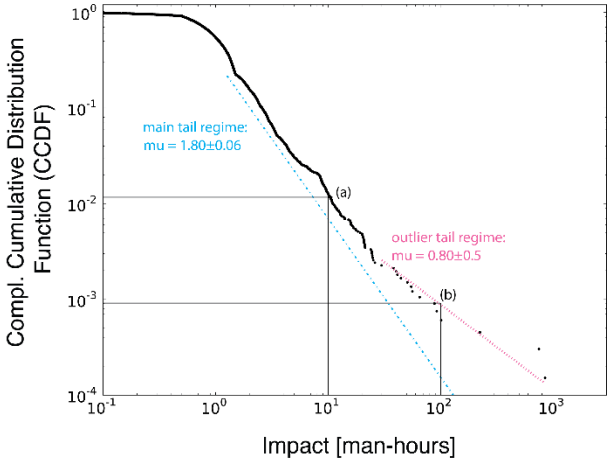

Fig. 3. The CCDF of investigation time for all incidents in the 11th time period. Note the three regimes: small incidents (here for impact < 0.5 man-hours), the main tail (blue dotted line), and an outlier tail (magenta dotted line).

The tail distribution shown in Figure 3 is best described by a main power law tail $P(C \geq c) \sim 1/c^{\mu}$, 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 because 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.

5.3 **Changes over time**

Cyber security and cyber threats are rapidly evolving, with new vulnerabilities announced on a daily basis. Over the six years included in our data, changes in security safeguards, network structure, and security processes have occurred. However, the impact distribution is consistently well accounted for by a power law model (see Table 2). Furthermore, we find evidence that all incidents taken together are overall becoming

less extreme over time. Figure 5 shows the impact distribution of all events greater than 2 man-hours over the time periods from 1 to 12. The exponent μ governing the skewness of the tail distribution of the power law for incidents' severity 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.

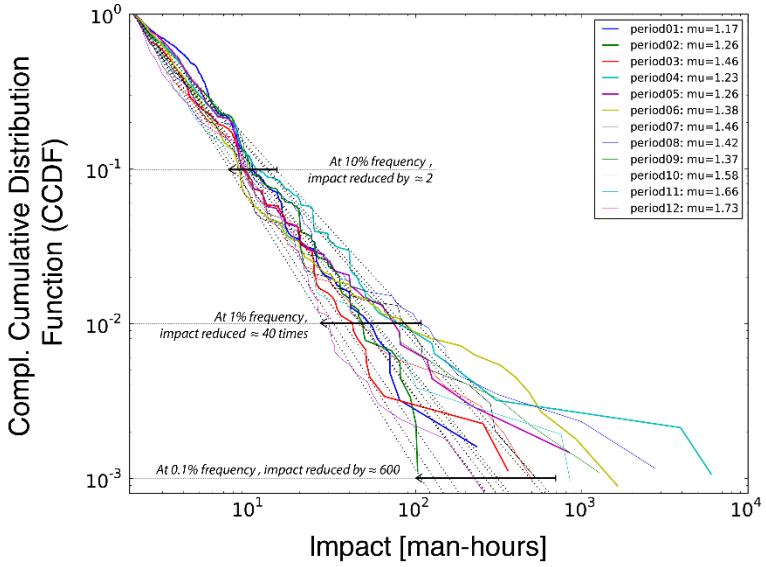

Fig. 4. Evolution of the complementary cumulative distribution function **(**CCDF) of impact for all incident types.

The CCDF represented in Figure 4 shows that the impact 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), the impact 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 reported in Table 2.

While the main tail section of the distribution (greater than 2 man-hours) has become less extreme over time, the number of small incidents has increased by a factor greater than 38 in period 12 compared to period 1.[13] Most of the rapid increase of the number of documented incidents over time was caused by the organization's implementation of new incident recording guidelines. The majority of additional incidents were small, requiring less than two hours of investigation. Table 2 shows the number of incidents during each of the 12 time periods observed from 2009 to 2014.

Table 2. 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 (**).

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
1	1,112	489	623	4,588	493	4,095	2	233	1.17	0.91	0.08	-
2	2,581	1,673	908	7,451	1,751	5,700	2	103	1.27	1.00	0.07	-
3	3,902	3,005	897	8,391	3,123	5,268	2	358	1.46	0.38	0.07	-
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
4*	3,455	2,535	920	9,326	2,596	6,729	2	307	1.24	0.83	0.07	
5	1,714	1,026	688	6,257	1,025	5,231	2	842	1.27	0.26	0.08	-
6	2,716	1,592	1,124	12,130	1,591	10,540	2	1,640	1.40	0.09	0.06	-
7	2,291	1,305	986	8,126	1,433	6,694	2	525	1.45	0.26	0.06	-
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)
8*	3,341	2,501	840	6,490	2,260	4,229	2	40	1.53	1.00	0.07	-
8**	3,355	3,341	14	11,774	6,490	5,284	41	2,762	0.72	0.50	0.53	-
9	3,774	2,879	895	10,317	2,840	7,478	2	1,267	1.34	0.14	0.07	-
10	5,922	4,892	1,030	11,017	4,685	6,332	2	501	1.59	0.15	0.06	-
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)
11*	6,618	5,595	1,023	9,799	5,457	4,342	2	30	1.80	1.00	0.06	-
11**	6,632	6,618	14	12,298	9,799	2,500	30	850	0.89	0.53	0.53	-
12	20,192	18,545	1,647	26,961	18,594	8,366	2	313	1.72	0.25	0.05	-

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 (i.e., 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 extreme event alone accounts for 2,762 hours (i.e., 115 days) of work. However, the majority of investigation time switches to small impact incidents in time periods 11 and 12. This shift corresponds to the new incident reporting guidelines described earlier. Interestingly, the increased diligence at low-level incidents appears to have only a small effect on the skewness of the tail distribution, but at a cost of a large increase

[13] Note, however, that most of the increase occurs from period 11 to 12, which corresponds to a directive to begin recording port scans and attempted website attacks.

in the total hours spent on the investigation of incidents. Figure 5B exhibits the linear relationship, with a slope = 0.28 ± 0.05 (p-value < 0.01), between the main tail exponent and the total number of incidents recorded during a time period.

Evidence also exists that incidents are generally becoming less extreme over time. Figure 5A shows that the exponent of the tail power law distribution for all incidents, as well as for each attack vector over 12 time periods, is increasing.

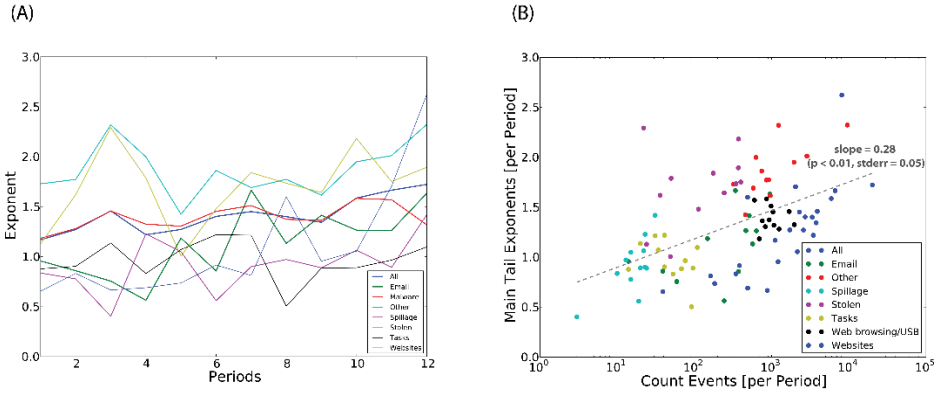

Fig. 5. A: Evolution of the power law exponent over time for different types of incidents; B: The main tail distribution (xmin = 2), characterized by the exponent μ, as a function of the number of events (logarithmic scale) per time period. The more events in one period, the larger μ, thus the less skewed the distribution. For a good understanding of the probabilities at stake, consider three values of μ = {0.5, 1, 2}. The probability of an incident of cost C is given by $C = 1/C^{\mu}$, so occurrence probabilities for events of cost 100 man-hours are respectively 10%, 1%, 0.01%.

5.4 Case study: Deployment of full-disk encryption (FDE) policy

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 in cases where the stored information is 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, μ = 0.52, p = 0.46). Comparing the distribution of severity before FDE policy implementation without the outliers (purple line, μ = 1.89, p = 0.79), and the distribution post-FDE implementation (blue line, μ = 1.85, p = 0.43), one cannot reject the hypothesis that they are both drawn from the same distribution. In other words, we can confidently say 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).

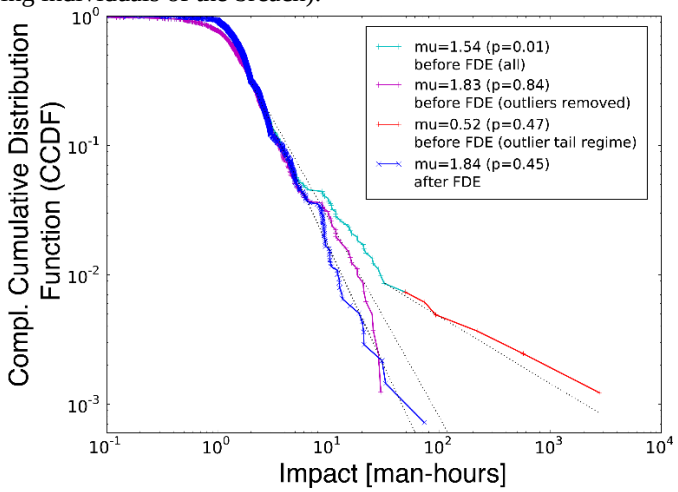

Fig. 6. The CCDF for lost devices before and after FDE was implemented.

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 (simply due to the removal of outliers), meaning that full disk encryption in that case has been a very cost-effective prevention measure.

6 Discussion

We have analyzed 60,000 incidents recorded at a large organization over a six-year period. While we find several changes in the rate of security incidents over time, many of the changes can be attributed to changes in the incident reporting guidelines. Therefore, the rate of many incident categories can be considered relatively constant over time. We also find that incidents fall into three categories of severity; namely small incidents, a main power law tail, and an outlier tail following another, more extreme power law. Incidents themselves are generally becoming less extreme over time as well,

indicating that the organization may be improving their detection and remediation. Finally, we empirically observe that the outlier tail distribution of investigation times for lost devices was removed after full disk encryption was implemented, meaning that the reported parameters can be used for further cyber risk modeling [6].

These results have a number of implications for decision and policy makers, including a demonstration of how incident management systems provide practical, operationally impactful insights about workflow and IT budgeting.

6.1 **Stability of cyber incident statistics**

At security conferences, a common concern is that a risk analyst will spend a significant amount of time calculating how often Windows 8 is attacked and exploited, but that the conclusions will not be relevant to Windows 10. In fact, the rate of cyber security incidents is constant, or changes on the timescale of years. It may be surprising that the rate of incidents is so regular given the number of changes that occurred in that organization over six years, which included employee turnover, network restructuring, adversary adaptations, evolving defensive capabilities, and leadership changes. Certain incident types (e.g. lost laptops and data spillage) are unlikely to disappear, but quantifying the rate of these predictable incidents can help in budget planning and resource allocation.

6.2 **Heavy-tailed Distributions**

There is considerable evidence that the impact of cyber incidents follows a heavy-tailed distribution [7,10,11]. These distributions are poorly characterized by their first and second moments (mean and variance). Failure to adequately address the heavy-tail nature of cyber attacks can deteriorate the workflow. On the contrary, excess risk could be transferred via an insurance product that provides coverage for rare, severe events that may overwhelm a SOC.

Incident's severity appears to be slowly decreasing over time, but the form of the tail distribution (power law) is consistent across categories and over time. It is unclear if the heavy-tailed distribution of event severity is a result of incident investigation, adversary skill, network topology, or some other cause. The decrease in the tail exponent over time could result from better investigation, better detection capabilities, or better response capabilities. It is important to note that many attack vectors (including website incidents, lost devices, and other incidents) have distinct outlier tails, which suggest a different growth process for some large incidents.

6.3 **Implications**

One of the most striking features of these data is how well they are modeled by relatively simple[14] and slowly evolving statistical models. Our results may be useful to safety engineers who develop cyber risk models [6]. It may also help managers test and implement new prioritization strategies [14,40,47]. While human judgment is critical for detecting and coping with risks, we believe that data-informed cyber risk models are also desirable to reduce some human biases in interpretation, which in turn result in suboptimal decisions [48].

Here, we stress that data about cyber security incidents have the potential to dramatically improve situational awareness by providing ground truths about the cost of cyber incidents. For example, another organization discovered that a significant portion of website attacks occurred against legacy servers that had been abandoned by a project team. Websites were not maintained or removed after the project expired, leading to severe breaches of security. Another Chief Information Security Officer (CISO) was able to show that contrary to the intuition of the company, the rate of attacks did not increase when the organization received media attention, and that additional staffing in anticipation of other attacks was not useful.[15]

7 **Limitations and Future Work**

As with any real world data set, incomplete records and recording biases exist. For example, security officers at the large organization studied here have reported that investigations taking less than 30 minutes are typically not documented and so we have an incomplete view of the smallest incidents. Also, these data do not address reputation damage, although qualitative evidence from the organization is consistent with literature that shows a tenuous link between data breaches and stock prices [49-51]. These data are limited to incidents that were detected. While it is possible that hackers have compromised this organization many other times but were not detected, challenging our results would require that the records used here are only a fraction of all attacks (detected and undetected). We believe that this is unlikely and that the results presented here are representative of the most common attacks that this organization has faced. Our work also says nothing about other organizations, although we hope that more incident data will become available to researchers in the future. Additional data

[14] Note that both the frequency of events described by a rate of event, and the severity described by a power law exponent, are 1-parameter statistics.

[15] Both examples come from private conversations with a large organization.

from a wide range of industries could greatly improve the awareness of different attacker strategies, and inform effective defenses.

This paper is a first attempt to understand how cyber security incidents are dealt with within organizations. More work on understanding the organizational challenges for cyber security is needed. In particular, strategies for managing heavy-tailed incidents and resource planning deserve more attention. For example, a comparison of a policy prioritizing the investigation of small incidents over large incidents (to test whether investigators become distracted when investigating large incidents) would benefit incident resource planning.

Our analysis presents a picture at a coarse-grained level, but the finer details and dynamics of cyber incidents are unknown. The waiting times between incidents and the duration of an incident investigation (in complement to the workload studied here) are also unknown. Other work could model the variability of incidents to determine if many small incidents precede larger incidents, therefore acting as an indicator of future attacks. A number of other unknowns remain, including how management decisions about investigation and remediation are made. Cyber incidents can track the effectiveness of new safeguards over time, but cannot necessarily predict which additional safeguards will be effective, especially in the case of adaptive adversaries.

8 Conclusion

Organizations are constantly bombarded with new cyber attacks, and security operation centers must continuously optimize their operations due to scarce resources. From tens of thousands of cyber security incidents, surprising trends emerge that can improve the cyber situational awareness of an organization, its resource investments, and its assessment of cyber risk. We find that the rate of cyber security incidents is relatively stable over time. Small incidents that are quickly resolved combine with a main power law tail and a transient extreme outlier tail regime to form impact distributions that are decreasing in incident severity over time and following security policy upgrades.

Acknowledgements. The authors would like to thank Alex Keller for help in this project, and Didier Sornette for helpful comments.

References

1. The White House, Office of the Press Secretary. “FACT SHEET: White House Summit on Cybersecurity and Consumer Protection. The White House.” February 13, 2015. Retrieved

from https://www.whitehouse.gov/the-press-office/2015/02/13/fact-sheet-white-house-summit-cybersecurity-and-consumer-protection

2. Department of Homeland Security, “Enhancing Resilience Through Cyber Incident Data Sharing and Analysis: Establishing Community-Relevant Data Categories in Support of a Cyber Incident Data Repository,” (September 2015), available at: www.dhs.gov/cybersecurity-insurance.
3. Florêncio, Dinei, and Cormac Herley. "Sex, lies and cyber-crime surveys. “Economics of information security and privacy III. Springer New York, 2013. 35-53.
4. Tøndel, Inger Anne, Maria B. Line, and Martin Gilje Jaatun. "Information security incident management: Current practice as reported in the literature." Computers & Security 45 (2014): 42-57.
5. Arora, Ashish, et al. "Measuring the risk-based value of IT security solutions." IT professional 6.6 (2004): 35-42.
6. Kuypers, M., and Pate-Cornell, E., "Quantitative Cyber Risk," Society for Risk Analysis Annual Meeting. Arlington, Virginia. December 7-9, 2015.
7. Wheatley, Spencer, Thomas Maillart, and Didier Sornette. "The Extreme Risk of Personal Data Breaches & The Erosion of Privacy." The European Physical Journal B 89.1 (2016): 1-12.
8. Cezar, Asunur, Huseyin Cavusoglu, and Srinivasan Raghunathan. "Outsourcing information security: Contracting issues and security implications." Management Science 60.3 (2013): 638-657.
9. Bolot, Jean, and Marc Lelarge. "Cyber Insurance as an Incentive for Internet Security." Managing information risk and the economics of security. Springer US, 2009. 269-290.
10. Maillart, T., and D. Sornette. "Heavy-tailed distribution of cyber-risks." The European Physical Journal B 75.3 (2010): 357-364.
11. Edwards, Benjamin, Steven Hofmeyr, and Stephanie Forrest. "Hype and Heavy Tails: A Closer Look at Data Breaches. “Workshop on Economics of Information Security”
12. Sornette, Didier, and Guy Ouillon. "Dragon-kings: mechanisms, statistical methods and empirical evidence." The European Physical Journal Special Topics 205.1 (2012): 1-26.
13. Schneier, Bruce. “Carry On: Sound Advice from Schneier on Security”. Wiley. 2013.
14. Anderson, Ross, and Tyler Moore. "The economics of information security." Science 314.5799 (2006): 610-613.
15. Frei, Stefan, et al. "Modeling the security ecosystem-the dynamics of (in) security." Economics of Information Security and Privacy. Springer US, 2010. 79-106.
16. Frei, Stefan, et al. "Large-scale vulnerability analysis." Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense. ACM, 2006.
17. Ralston, Patricia AS, James H. Graham, and Jefferey L. Hieb. "Cyber security risk assessment for SCADA and DCS networks." ISA transactions 46.4 (2007): 583-594.
18. Tidwell, Terry, et al. "Modeling internet attacks." Proceedings of the 2001 IEEE Workshop on Information Assurance and security. Vol. 59. 2001.

19. Liu, Debin, XiaoFeng Wang, and L. Jean Camp. "Mitigating inadvertent insider threats with incentives." Financial Cryptography and Data Security. Springer Berlin Heidelberg, 2009. 1-16.
20. Liu, Debin, XiaoFeng Wang, and Jean Camp. "Game-theoretic modeling and analysis of insider threats." International Journal of Critical Infrastructure Protection 1 (2008): 75-80.
21. Lacey, David. Managing the Human Factor in Information Security: How to win over staff and influence business managers. John Wiley & Sons, 2011.
22. Kjaerland, Maria. "A taxonomy and comparison of computer security incidents from the commercial and government sectors." Computers & Security 25.7 (2006): 522-538.
23. Arora, Ashish, Rahul Telang, and Hao Xu. "Optimal policy for software vulnerability disclosure." Management Science 54.4 (2008): 642-656.
24. Arora, Ashish, and Rahul Telang. "Economics of software vulnerability disclosure." IEEE security & privacy 1 (2005): 20-25.
25. Zhang, Su, Doina Caragea, and Xinming Ou. "An empirical study on using the national vulnerability database to predict software vulnerabilities." Database and Expert Systems Applications. Springer Berlin Heidelberg, 2011.
26. Neuhaus, Stephan, et al. "Predicting vulnerable software components." Proceedings of the 14th ACM conference on Computer and communications security. ACM, 2007.
27. Shin, Yonghee, et al. "Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities." Software Engineering, IEEE Transactions on 37.6 (2011): 772-787.
28. D'Ambros, Marco, Michele Lanza, and Romain Robbes. "An extensive comparison of bug prediction approaches." Mining Software Repositories (MSR), 2010 7th IEEE Working Conference on. IEEE, 2010.
29. Rhodes, C. J., and Roy M. Anderson. "Epidemic thresholds and vaccination in a lattice model of disease spread." Theoretical Population Biology 52.2 (1997): 101-118.
30. Killcrece, Georgia, et al. State of the practice of computer security incident response teams (CSIRTs). No. CMU/SEI-2003-TR-001. CARNEGIE-MELLON UNIV PITTSBURGH PA SOFTWARE ENGINEERING INST, 2003.
31. Chew, E., Swanson, M., Stine, K., Bartol, N., Brown, A., Robinson, W. Performance Measurement Guide for Information Security. National Institute of Standards and Technology Special Publication 800-55 rev. 1. July 2008. http://csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf
32. Tipton, Harold F., and Micki Krause. Information security management handbook. CRC Press, 2003.
33. Kruse II, Warren G., and Jay G. Heiser. Computer forensics: incident response essentials. Pearson Education, 2001.
34. Ahmad, Atif, Sean B. Maynard, and Graeme Shanks. "A case analysis of information systems and security incident responses." International Journal of Information Management 35.6 (2015): 717-723.

35. Condon, Edward, Angela He, and Michel Cukier. "Analysis of computer security incident data using time series models." Software Reliability Engineering, 2008. ISSRE 2008. 19th International Symposium on. IEEE, 2008.
36. Miani, Rodrigo Sanches, et al. "An Empirical Study of Connections Between Measurements and Information Security."
37. Joh, HyunChul, and Yashwant K. Malaiya. "Seasonal variation in the vulnerability discovery process." Software Testing Verification and Validation, 2009. ICST'09. International Conference on. IEEE, 2009.
38. Bando, Koichi, and Kenji Tanaka. "Trend Analyses of Accidents and Dependability Improvement in Financial Information Systems." Dependable Computing (PRDC), 2011 IEEE 17th Pacific Rim International Symposium on. IEEE, 2011.
39. Böhme, Rainer. "Security metrics and security investment models." Advances in Information and Computer Security. Springer Berlin Heidelberg, 2010. 10-24.
40. Jerman-Blažič, Borka. "Towards a standard approach for quantifying an ICT security investment." Computer Standards & Interfaces 30.4 (2008): 216-222.
41. Thomas, R. C., Antkiewicz, M., Florer, P., Widup, S., & Woodyard, M. (2013). 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).
42. Clauset, Aaron, Cosma Rohilla Shalizi, and Mark EJ Newman. "Power-law distributions in empirical data." SIAM review 51.4 (2009): 661-703.Embrechts, P., Resnick, S. I. & Samorodnitsky, G. Extreme value theory as a risk management tool. North American Actuarial Journal 3, 30–41 (1999).
43. Embrechts, P., Resnick, S. I. & Samorodnitsky, G. Extreme value theory as a risk management tool. North American Actuarial Journal 3, 30–41 (1999).
44. Embrechts, P., Kluppelberg, C. & Mikosch, T. Modelling extremal events: for insurance and finance. Vol. 33. Springer (1997).
45. Stephens, M. A. EDF statistics for goodness of fit and some comparisons. Journal of the American statistical Association 69.347 (1974).
46. Wheatley, S. and Sornette, D. Multiple Outlier Detection in Samples with Exponential and Pareto Tails: Redeeming the Inward Approach and Detecting Dragon Kings, Journal of Applied Statistics (in press). (http://arxiv.org/abs/1507.08689 and http://ssrn.com/abstract=2645709).
47. Maillart, Thomas, et al. "Quantification of deviations from rationality with heavy tails in human dynamics." Physical Review E 83.5 (2011): 056101.
48. Mersinas, Konstantinos, et al. "Experimental Elicitation of Risk Behaviour amongst Information Security Professionals." WEIS 2015.
49. Campbell, K., Gordon, L., Loeb, M., and Zhou, L. "The economic cost of publicly announced information security breaches: empirical evidence from the stock market," Journal of Computer Security (11:3) 2003, p 431.

50. Cavusoglu, H., Mishra, B., and Raghunathan, S. "The effect of Internet security breach announcements on market value: Capital market reactions for breached firms and Internet security developers," International Journal of Electronic Commerce (9:1), Fall 2004, pp 69-104.
51. Kannan, K., Rees, J., and Sridhar, S. "Market reactions to information security breach announcements: An empirical analysis," International Journal of Electronic Commerce (12:1), Fall 2007, pp 69-91.