Analysis of the impact of cyber events for cyber insurance

[1], Kjartan Palsson[1], Steinn Gudmundsson[1], and Sachin Shetty*[2]

[1]Department of Computer Science, School of Engineering and Natural Sciences, University of Iceland
[2]Virginia Modeling, Analysis and Simulation Center, Old Dominion University, 2026 Constant Hall, Norfolk, VA 23529, USA

Abstract

The mass adoption of cyber insurance will be predicated on the ability to conduct quantitative cyber risk assessment. This capability is crucial for not only providing insights into the cost of targeted threats but also provide incentives for insured enterprises to invest in protection aimed at preventing exploitation of targeted threats. Research indicates that asymmetric information, correlated loss, and interdependent security issues make this difficult if insurers cannot monitor the cybersecurity efforts of the insured enterprises. In this paper, we present an analysis of cyber impacts based on cyber incidents reported in the Advisen cyber loss data feed. We show i) how exposure to cyber incidents varies between corporate sectors; ii) how the type of an incident relates to the number of entities and individuals affected by it; iii) how the type of incident relates to the eventual financial cost; iv) what type of information is most frequently compromised; v) breakdown of the main actors behind cyber incidents; and vi) how tree-based classifiers can be used to gain insights into cyber risk indicators affecting the cost of incidents.

*Corresponding Author: sshetty@odu.edu

1 Introduction

Cyber risk assessment is paramount for organizations to conduct effective risk management and enhance resilience to cyber threats (Refsdal et al. 2015). Quantitative cyber risk assessment also benefits insurers to provide affordable and comprehensive insurance coverage (SwissRe 2017). However, current cyber risk assessment processes are primarily qualitative and lack quantitative insights. In order to conduct effective cyber risk assessment, it is critical to quantitatively estimate the frequency of incidents and the severity of potential losses across sectors.

The processes of distinguishing insureds based on cyber risk levels is a challenging task for insurers and underwriters. Currently, in order to conduct cyber risk assessment, the insurer often requires an organization to undergo an application process and several underwriting meetings. In addition, it is challenging for insurers to differentiate risks at first glance and they therefore charge similar premiums for organizations across sectors and within same sector. The state-of-the practice solution of providing rebates to the policyholders that have good historical cybersecurity posture is not tenable. Though, this practice can keep the risks endured by insurers manageable, but the lack of convenient and proactive risk assessment measures makes cyber insurance less than satisfactory to buyers, leading to limited coverage. Cyber insurance products are often criticized for their high premium rates, obscure policy languages and limited capacities (RMS 2019; Betterley 2013; Romanosky et al. 2019; Orlando et al. 2017; Eling and Schnell 2016).

In order to conduct quantitative cyber risk assessment, there is a need to develop the ability to analyze historical cyber incidents from verifiable sources and perform empirical analysis to identify factors that determine their frequency and severity. In addition to analysis, there is also a need for statistical models coupled with these factors as parameters for predicting future incident frequency and loss severity. To extend our understanding of the losses resulting from cyber incidents, there is a need to include both cyber incident data and corporate financial data so that we can study how cyber incidents impact the financial

performance of affected companies.

Jacobs et al. (2016) developed a Cyber Value at Risk model that they applied to approx. 50 Dutch companies using data from financial statements, company income, assets and IP. The authors estimated that the yearly loss for the Dutch economy is approximately 10 billion euros. The cyber risk exposure was highest in the banking, defense and aerospace, technology and electronics, and public sectors. Lagazio et al. (2014) deployed a system dynamics model to analyze the impact of cyber crime on the financial sector. The results showed that emphasis on protecting customer trust and loyalty and over-spending on defence measures are important factors in the cost cyber crime. The author also found that cyber incidents were chronically under-reported. Dreyer et al. (2018) describe a methodology for estimating present and future financial cost of cyber risk. The methodology enables identification of the value at risk by country and industry sector; computes direct and indirect costs based on data from OECD. Analysis of three case studies demonstrated that the cost estimates were highly sensitive to modeling assumptions, different sets of plausible parameter values resulted in cost estimates differing by a factor of 20.

Although there are several products on the market that aim to tackle the problem of cyber risk assessment, we find that they have several limitations. The products tend to focus on exploitability of vulnerabilities rather than quantifying risk to critical cyber assets, They provide limited information about internal cost activity centers such as detection, investigation, containment and recovery, as well as external costs due to information loss, business disruption, equipment damage and revenue loss. The available products provide little insight into how to balance remediation ROI and total cost and more generally they do not describe the impact of cyber attacks on the business processes of organizations. Tools such as, Bitsight and SecurityScorecard provide cyber risk scores, the scores are not correlated with financial losses (Shetty et al. 2018). First of all, computing premiums on cyber insurance polices rely largely on expert judgement, rather than quantitative data. In addition, few of the tools evaluate cyber risks from a monetary perspective. Last but not

least, because cyber incident data is difficult to collect, many products and services rely on obsolete datasets to assess cyber risks, which may fail to capture the ever changing risk landscape (M. Eling and J. H. Wirfs 2015; Martin Eling and J. Wirfs 2019a).

Though there are several challenges in realizing an effective cyber insurance policy, shortage of data plays a significant role. There is a lack of incentives for companies to share cyber threat information in a public or permissioned repository. To address this issue, parameterized models that mimic cyber behavior have been developed. For instance, Maillart and Sornette (2010), Wheatley et al. (2016), Edwards et al. (2016), and Martin Eling and J. Wirfs (2019b) have indicated that risk from certain cyber events are characterized by heavy-tailed distributions. However, this approach has been found to be somewhat limited (Schnell 2018).

There have been some efforts recently to develop data driven cyber risk assessments and evaluate their implications on cyber insurance. Biener et al. (2015) investigated the inadequacy of insurance for managing cyber risk. Based on 994 cases of cyber losses extracted from an operational risk database, they investigated the insurability of cyber risk by systematically reviewing the set of criteria that classify risks in terms of actuarial, market, and societal conditions and were able to bring to light highly interrelated losses, lack of data and severe information asymmetries. Kunreuther and Pauly (2018) describe results of respondents stated preferences for purchasing insurance for low-probability, high-consequence events where the probability of a loss and its consequences are stable over time. Pooser et al. (2018) used data from publicly traded U.S. property–casualty insurers, to study cyber risk identification from 2006 to 2015 and firms' characteristics related to cyber risk perception. Ashby et al. (2018) showed an instance of how German banks manage cyber risks based on interviewing IT managers, risk managers and external experts from ten participating banks. Smidt and Botzen (2018) presented analysis of individual perceptions of cyber risks based on data collected from survey of corporate professionals.

Romanosky (2016) examined 12,000 cyber events in the Advisen cyber loss data feed

(Advisen 2019), characterizing cyber breaches by industry and identifying those industries that were affected the most. Romanosky correlated costs to bad debts and fraud within other industries to look for impacts of cyber events on the financial outlook of a company. The main finding was that the cost of a typical cyber incident was less than $200k and accounts for only 0.4% of the organization's estimated annual revenues.

In this paper we present an analysis on cyber loss events and their financial impacts, using Advisen's cyber loss data feed. We discuss how our results can be used to develop effective cyber insurance policies. Our models provide coarse grained predictions based on discretized loss and realize the following benefits: i) a data-driven analysis of cyber loss events which provides insight into how cyber loss varies with sector type and firm size; and ii) modeling of the financial cost of different types of cyber incidents that can be used to develop informed cyber insurance policies.

2 Methodology

In this section, we provide a detailed description of the cyber loss data feed provided by Advisen. The cyber loss feed has facilitated the development of a model that can be used to gain insight into the importance of factors that impact cost of cyber incident. We also provide a high-level overview of Random Forests classifiers, which is the modeling methodology applied to the Advisen dataset.

2.1 Data Description

We have acquired the cyber loss data feed from Advisen, a US-based organization that acquires and sells cyber loss and incident data to insurers, reinsurers, brokers and cyber modeling firms (Advisen 2019). The cyber loss data feed was generated by processing cyber loss events reported by government websites (Securities & Exchance Commission (SEC), Federal Trade Commission (FTC), Federal Communications Commission (FCC), US De-

partment of Homeland Security (DHS), State data breach of notification websites, State Freedom of Information Act (FOIA) requests) and other sources, that include, official court and litigation sources, and websites focusing on data breaches, data security vendors, annual reports, and analysis (Advisen 2019). The total number of incidents in the data set analyzed here is approximately 75,000 involving 45,000 companies all over the world. US companies account for 86% of the companies listed. The analysis in this paper is based on all the incidents, unless otherwise stated. There are 16 types of cyber incidents, distributed across 20 corporate sectors, with most of the incidents taking place within the last decade. The incidents include highly publicized cases such as Enron, Vioxx and major Yahoo data breaches, as well as many incidents involving single individuals. Cyber incidents are associated with a parent firm and include information such as, case type, case status, affected count, event date, source of loss, type of loss, type of actor, loss amount, company size, company type, number of employees and the outcome of an eventual court case when applicable. The loss associated with a malicious data breach event is calculated by including fines levied by FINRA, response costs and legal fees associated with lawsuits for settlement.

An entry in the database represents an incident that occurred at a firm which has resulted in a potential significant financial loss that is not guaranteed to be insured. For most of the entries, the actual financial loss is known, however, there are instances where non-monetary information is reported. In situations, when there are multiple entities associated with a firm, the case is associated with the parent firm and the cumulative impact is attached to only one entity. However, in cases where there is known penalty or losses that impact more than one entity within the organization, there are separate cases assigned to each entity.

Figure 1 shows the number of cyber incidents for the five most common incident types in the period 2000 to 2018. The large drop in 2017 and 2018 is in part explained by delays in registration of incidents into Advisen's database. The difference between the date of incident and the date when it is entered into the database ranges from 2 to 5 years on average, depending on incident type. The figure shows that since 2005 there has been a

large increase in incidents involving malicious data breach as well as incidents involving unauthorized contact or disclosure (C/D). The number of incidents involving unintentional disclosure of data is also on the rise while the number of incidents involving physically lost or stolen data has been fairly steady in the period. Typically, cyber events go unreported as organizations do not have a clear incentive to share cyber threat information. Thus, it is challenging to quantify how close the incident numbers in the database are to the total number of cyber incidents. Some registration bias is likely to be present in the database, e.g. the introduction of a new data source by the database maintainers could lead to a disproportionate increase in a particular type of incident.

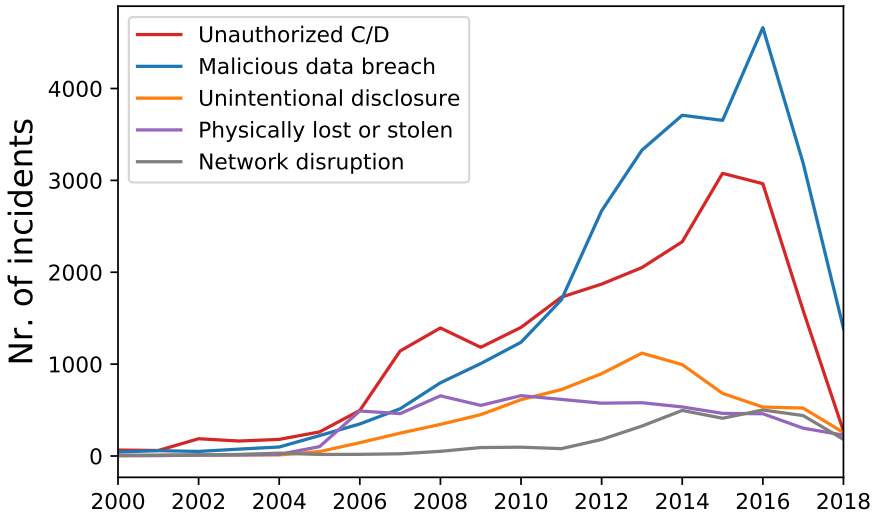

Figure 1: Number of cyber incidents in the Advisen cyber loss data feed, occurring between 2000 and 2018.

2.2 Random Forests classifiers

Random Forests is a classification algorithm based on decision trees (Breiman 2001) which has been found to have good prediction accuracy on many real-world problems, e.g. prediction of fraudulent auto insurance claims (Derrig and Francis 2006), insurance purchases (Wu

et al. 2017), customer retention (Guelman et al. 2012) and more. The input to the algorithm is a set of n input – output pairs (called the training set), $\{(x_i, y_i)\}_{i=1}^{n}$ where the inputs x_i are real valued vectors, e.g. features of a particular cyber incident, and the outputs y_i take on discrete values corresponding to the different classes, e.g. low, medium and high financial cost. The goal is to learn a rule which predicts the class of an unseen example x with high accuracy.

Random Forests is an ensemble method where multiple decision trees are constructed from the training set by bootstrap sampling. To predict the class of an example x, it is sent down all the trees. Each tree predicts a class and the final prediction is determined by majority voting amongst all the trees. Random Forests can be used to obtain a measure of the importance of individual inputs (a.k.a. features) with respect to classification accuracy as follows. Since each each decision tree is fitted to a bootstrap sample of the original data set, the left out examples, the so-called out-of-bag examples, can be used to estimate error rates for individual trees. Furthermore, by randomly permuting feature values of the out-of-bag examples, one feature at a time, and sending the permuted examples down the tree again, the increase in error rate provides a measure of the importance of the feature. This methodology, referred to as permutation-based feature importance, is used below. Discrete features were encoded using one-hot encoding where a k-valued discrete feature is replaced by k binary features where binary feature i is 1 if and only if the feature value is equal to i.

In the experiments that we conducted, classifier accuracy was estimated by randomly splitting the data set into disjoint training (80%) and test (20%) sets and computing the fraction of correctly predicted test cases. This was repeated 100 times and the results averaged[1].

1. The Random Forests implementation in scikit-learn (Pedregosa et al. 2011), version 0.21.3, was used with the default settings, except the number of trees was set to 200. The `permutation_importance` score in scikit-learn was used to obtain feature importance scores.

3 Results

In this section, we provide detailed presentation of the results based on statistical analysis and application of a Random Forests classifier. The statistical analysis provides insights into the frequency and type of cyber incidents across various sectors and type of perpetrators. The cyber risk modeling using a Random Forests classifier provides indications on the number of affected entities and factors that impact cost of cyber incident.

3.1 Sector comparison

Figure 2 shows the total number of incidents for each of the 20 corporate sectors included in the data set (bar chart). The figure shows clearly that "malicious breach" incidents are prominent across all sectors. There are many cases labeled as "unauthorized contact/disclosure" but they are mainly concentrated in the finance sector as well as the administrative, support, waste management and remediation (ASWR) services sector. The sectors differ considerably in size and this needs to be taken into account when comparing incident rates between sectors. Using recent data for the number of businesses in US sectors (NAICS 2019), the figure also includes the incident rate for US firms (black dots).

The *incident rate* for sector i is defined as the ratio of the number of US incidents in sector i to the number of US establishments in sector i, i.e. it is the relative frequency of incidents in sector i. Figure 2 shows how the incident rate varies across sectors. The accommodation and food services as well as ASWR services have the highest incident rates, indicating that these sectors are more susceptible to cyber incidents than other sectors. On the other hand, the finance, information, health care, public administration, retail and wholesale sectors all have similar incident rates. This suggests that their exposure to cyber events is similar, although the type of incidents differ. Unintentional disclosure of data and physically lost/stolen data account for a large proportion of incidents in the health care and public administration sectors while malicious breach and unauthorized C/D are more

prominent in the other four sectors.

The distribution of cyber incidents across the sectors in Figure 2 provides insights to companies on strategies for self-insurance vs. cyber insurance. By analyzing the frequency of occurrence of particular cyber events, organizations can decide events that should be included in a cyber insurance policy vs. events that should be self-insured. Most organizations lack quantitative insights that would aid them in selecting the relevant cyber insurance policy that includes coverage for events that routinely affect their sector.

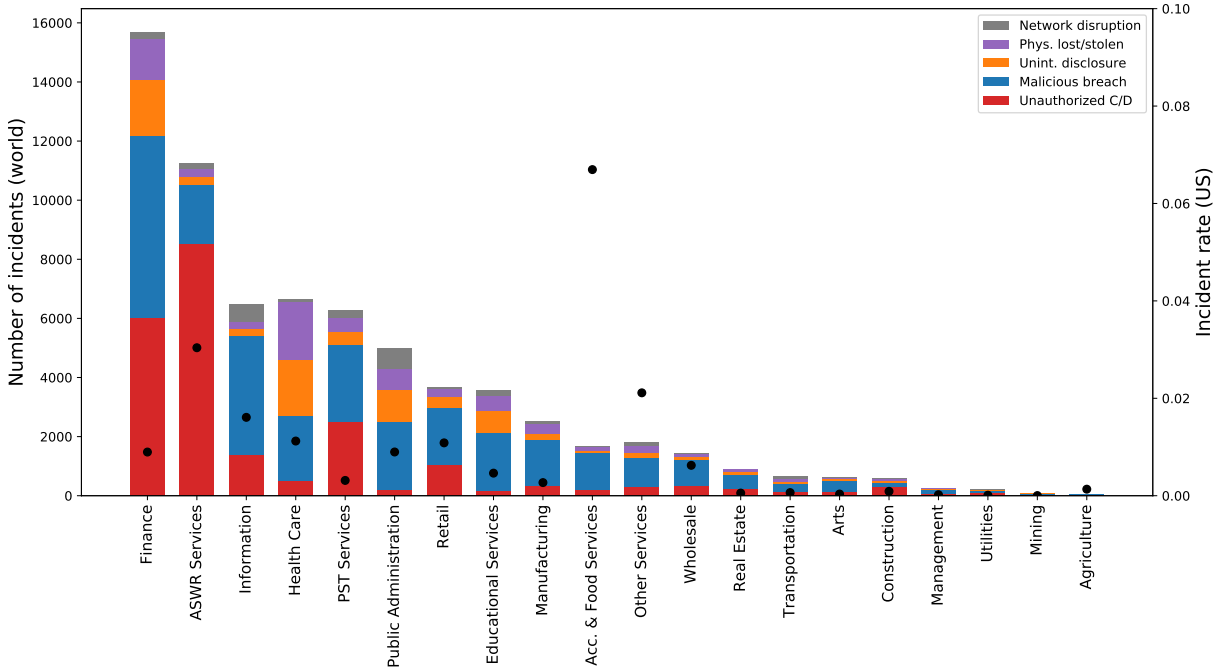

Figure 2: The number of cyber incidents in each of the 20 corporate sectors (bar chart, left axis) and the incident rate for each sector in the US (black dots, right axis). Abbreviations: Administrative, support, waste management and remediation (ASWR), Professional, scientific, and technical services (PST).

3.2 Types of incidents

Figure 3 summarizes the severity of the different types of incidents in the data set, both in terms of cost (left panel) and in terms of the number of affected individuals or entities (right

panel). Phishing[2] incidents which often target specific individuals or organizations turn out to be very costly, with the median cost around half a million dollars. This is followed by malicious breach incidents where the median cost is around one hundred thousand dollars, with some extreme outliers. On the other end of the spectrum are cyber extortion cases, exemplified by ransomware attacks which often target a large group of users indiscriminately. The median cost for these cases is only a few thousand dollars. Cases involving physically lost or stolen data and unintentional disclosure of data are also very costly but are of a somewhat different nature as discussed below.

The number of affected entities (identities breached or stolen, devices compromised etc.) follows a similar pattern as the total cost, in the sense that the most costly incident types tend to involve the largest number of entities. Most of the unauthorized contact/disclosure incidents involve only a single individual but there are also some extreme incidents in this category, estimated to have affected hundreds of millions of users.

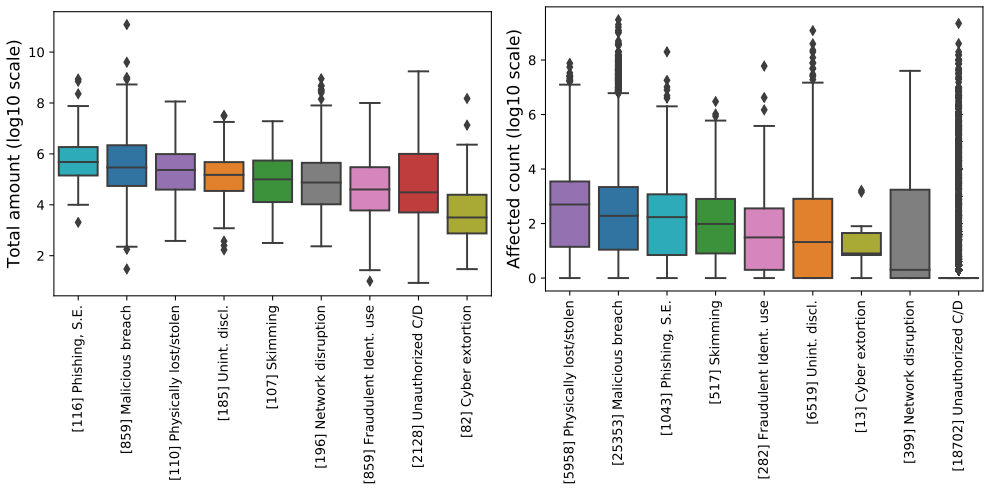

Figure 3: Severity of incidents measured in terms of total cost (left, log-scale) and affected count (right, log-scale). The boxes show the interquartile range, horizontal lines the median, whiskers indicate 1.5 times the interquartile range and dots represent outliers. The number of incidents in the Advisen data set for each type is given in brackets.

Figure 4 shows that the incidents leading to breach of data differ considerably for the three

2. Phishing refers to the practice of tricking people into revealing sensitive information such as username and passwords, bank account information etc., via email, social media or fake websites.

major types of information compromised, personal identity information, personal health information, and personal financial identity. Malicous data breaches feature prominently in all cases and are the cause of over half the cases involving financial information. Unintentional disclosure of data and data that is physically lost or stolen are the second and third leading causes for finance information breaches. Debit or credit card information is involved in 28% of all the financial information cases. Most of the cases involving financial information are attributed to "bad actors", highlighting the need for financial companies to employ strict cyber security measures. In contrast, malicious intent appears to be less of an issue in the case of health information breaches. Mishandling of data is a large factor and could likely be reduced by introducing appropriate data handling protocols and by employee training. In case of breach of personal identity information, unauthorized contact or disclosure (C/D) account for almost 60% of the cases. The leading incident types will now be described in more detail.

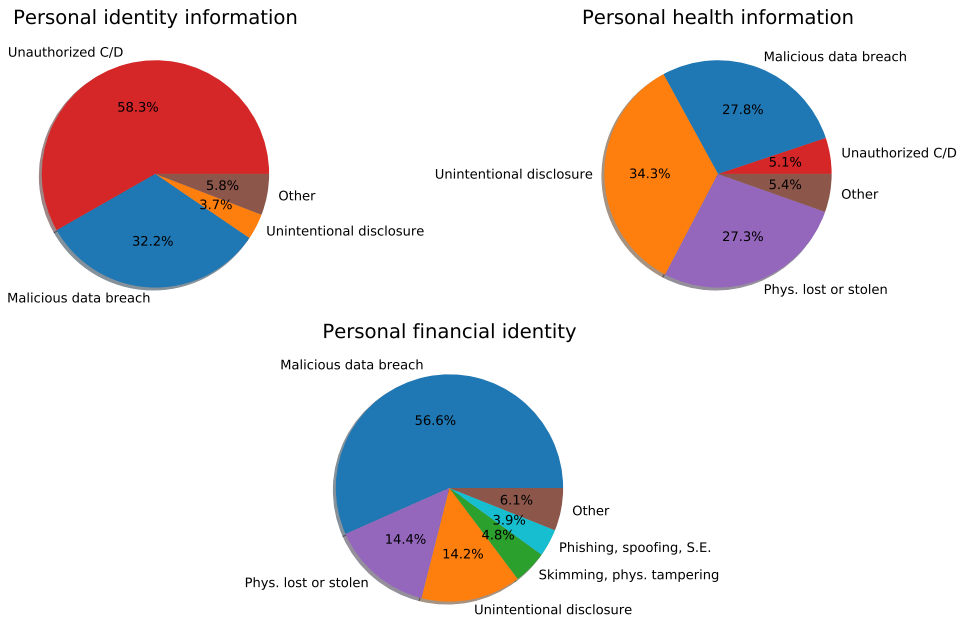

Figure 4: Information that is compromised in relation to the most common types of incidents.

Examination of incidents from the unauthorized contact or disclosure class reveals that although 80% of them involve only a single individual (see figure 3) the incidents are highly

heterogeneous. They range from disputes involving the use of photographs on underwear packaging to a Yahoo data breach involving 200 million users.

The specific laws violated are most often fair debt collection protection acts, telephone consumer protection acts and fair credit reporting acts. The incidents mostly involve shady debt collectors, telemarketers, robocalls and credit rating issues.

Malicious data breach is the most frequent form of cyber incidents across almost all corporate sectors. They affect more individual or entities than any other type of data breach and are among the most expensive incidents. Servers, websites and email are most often compromised and the data breached is mostly in the form of financial and personal information (over 90%). The responsible actors are mostly unknown, but approximately 7% of the instances are attributed to current or former employees and 2% to hacktivists[3].

Unintentional data disclosure represents the third most frequent type of data breach. Printed records are most often compromised, followed by email, servers and websites. Financial, health and personal information are almost exclusively exposed.

From a cyber insurance policy perspective, the breakdown of different types of malicious data breaches provides additional insights into the type of data breach coverage that needs to be included in the policy.

3.3 Perpetrators of cyber incidents

The Advisen data set contains information on the actors or perpetrators involved for some of the cyber incidents. Table 1 groups the most common perpetrators according to incident type and the type of information exposed. Cyber incidents attributed to hacktivists and foreign nation states have several things in common. They involve mostly malicious data breach and network disruption. The loss is mostly personal identity information and corporate business income/services. Incidents attributed to terrorists exhibit a similar pattern to the other two. A second cluster is formed by employees, vendors, consultants and trusted third parties,

3. A hacktivist tries to bring about political or societal change using hacking.

suggesting that these four actors share a similar risk profile.

Criminal organizations form a singleton cluster. They mostly target financial information via malicious data breach and network disruption. Other means include skimming, physical tampering and phishing (data not shown). These observations help in adding problematic exclusions in the cyber insurance policy. For instance, the insights on frequency of foreign criminal organizations being the source of attacks in a particular sector will help ensure that the policy doesn't exclude acts of foreign enemies.

Perpetrator		Cases	Malicious breach	Unintent. disclosure	Privacy breach	Phys. lost/ stolen	Network disruption	PII	PFI	PHI	CLDA	CLBIS
Crim. Org.	E	511	0.40	0.01	0.00	0.04	0.27	0.14	0.45	0.04	0.18	0.14
Hacktivist	E	1122	0.47	0.00	0.00	0.00	0.51	0.38	0.06	0.01	0.09	0.44
Nation State	E	436	0.39	0.00	0.00	0.01	0.46	0.25	0.07	0.03	0.15	0.43
Other	E	29883	0.70	0.02	0.00	0.14	0.04	0.34	0.51	0.08	0.01	0.04
Terrorist	E	200	0.26	0.01	0.02	0.02	0.64	0.32	0.00	0.00	0.06	0.60
Vendor	E	611	0.25	0.32	0.01	0.34	0.01	0.15	0.64	0.19	0.01	0.01
Consultant	I	119	0.35	0.30	0.01	0.24	0.03	0.17	0.52	0.26	0.03	0.03
Employee	I	4881	0.37	0.41	0.04	0.09	0.00	0.16	0.45	0.35	0.02	0.01
Organization	I	27387	0.02	0.13	0.79	0.01	0.00	0.83	0.09	0.05	0.00	0.01
3rd Party	I	673	0.29	0.38	0.00	0.12	0.06	0.14	0.51	0.25	0.02	0.09

Table 1: Actors or perpetrators behind the most common cyber incidents, classified as external (E) or internal (I). The fraction of incidents attributed to a given perpetrator is shown in green and the type of information exposed in purple. Abbreviations: Personal identity information (PII), personal health information (PHI), personal finance identity (PFI), Corporate loss of digital assets (CLDA) and Corporate loss of business income/services (CLBIS).

3.4 Cyber Risk Modeling

In current state-of-the practice cyber insurance policies, most organizations are categorized as high risk due to lack of cyber risk insights (Lu et al. 2018; Huth 2018; Kesan 2017). Though, it is very challenging to derive fine grained cyber insurance premium policies, the first step to achieving this goal should be to categorize organizations into finite categories of cyber risk, for example, low,medium and high.

It is of interest to classify the outcome of a cyber incident based on sector/organization information, type of incident, e.g. data breach or network disruption. For example, the impact of a cyber exploit that was used to commit data breach on a server hosting patient health records vs. network disruption that impacts website loading times is likely to be

different. Whether such models are successful or not depends on several factors. In a classification setting, the main determining factor is of course whether there is any inherent difference between the classes of interest, and whether the data set contains relevant inputs (features) that capture the difference. Other factors include the amount of data available and whether the model is sufficiently rich to capture the difference between classes. Data quality is yet another factor, e.g. missing values can play a large role.

We now discuss attempts at predictive modeling using the Advisen data set. In particular, we try to predict whether a resulting court case was settled or dismissed; the total cost of the incident to the company involved; and the number of individuals or entities affected by a cyber incident.

3.4.1 Determining settled vs. dismissed incidents

The Advisen data set contains several entries where a cyber incident is followed by a court case. If the outcome of a court case is known, this information is added to the data set. The most frequent outcomes are incidents that were either settled or dismissed in court[4]. The Advisen data set contains approximately 100 features for each incident but inspection of missing values shows that many features are only present for a tiny fraction of the incidents. We ended up using the following six features: 1) *incident type* 2) *data source* 3) *data type* 4) *number of employees* 5) *yearly company revenue* 6) *company sector*. Features 4 and 5 are continuous valued but the rest discrete. This resulted in a data set with 6,800 incidents that were labelled ”settled” and 4,800 incidents that were labelled ”dismissed”.

4. a third class present in the data, ”dismissed without prejudice.” ($n = 1,800$), is omitted here.

Feature	Settled	Dismissed
Incident type	Unauthorized C/D - 87.9% Malicious breach - 5.0%	Unauthorized C/D - 94.4% Malicious breach - 2.7%
Data type	PII - 91.0% PFI - 6.0% PHI - 2.5%	PII - 94.6% PFI - 4.2% PHI - 1.0%
Data "source"	Tel. comm. - 58.8% Privacy laws - 28.5% Server - 5.0%	Tel. comm. - 53.4% Privacy laws - 39.8% Server - 2.9%
Employees	Median: 152 IQR: 1630	Median: 200 IQR: 1773
Revenues (million USD)	Median: 18.6 IQR: 241.1	Median: 21.0 IQR: 240.4
Sector	ASWR services - 37.9% Finance - 25.4% PST services - 11.0%	ASWR services - 43.7% Finance - 25.1% PST services - 13.1%

Table 2: The most frequent feature values for the dismissed and settled incidents. The following abbreviations are used: Contact or disclosure (C/D), Personal Identity Information (PII), Personal Financial Identity (PFI) and Personal Health Information (PHI).

Table 2 provides a summary of the underlying feature values for the two classes. The table shows that no single feature is likely to clearly discriminate between the two. Therefore, we decided to build a Random Forests classifier where all six features were used as inputs, and the output is either "Settled" or "Dismissed". To avoid potential problems with class imbalance, we discarded 2000 "settled" instances at random to obtain a data set with equal number of instances in each class.

The accuracy of the resulting classifier was found to be 55.6% (SD=1.3%), i.e. only slightly better than chance. One possible explanation is that almost all the incidents are of the unauthorized contact or disclosure type. As discussed above, this class is extremely heterogeneous which is likely to cause difficulties in prediction. The above experiments suggest that in order to accurately predict the outcome of court cases, we need to involve additional features not found in the Advisen data set. This information might include cyber threat information (CISA) that would include details of the threat, specific vulnerabilities that were exploited, availability of patches, etc. For example, if an organization has been a

victim of a cyber attack resulting from a vulnerability for which a patch has been existing for some time, then cost of the cyber incident should be increased significantly. This was an fixable issue and is an indication of lack of cyber security controls in the organization.

3.4.2 Determining the number of affected entities

Another outcome of interest is the number of entities affected by a given cyber incident. Figure 3 shows that for most of the unauthorized C/D cases, only a single entity is affected. To simplify the prediction task, we decided to group the affected count into three bins, corresponding to a single entity, 2 – 1000 entities and more than 1000 entities (the value of 1000 was chosen somewhat arbitrarily). To avoid class imbalance, we discarded data from the two largest classes to end up with 6032 incidents in each class. We again fitted a Random Forests classifier, using the same features, classifier configuration and evaluation methodology as before. The resulting classifier had 60.1% accuracy (SD 1.1%) which is significantly better than random guessing (33% accuracy). Examination of prediction errors for each class showed that the accuracy for the single entity class was slightly higher than for the other two. Figure 5 (left panel) shows the features which had the biggest impact on the number of affected entities. As we can see, the incident type plays the largest role in predicting the outcome while the company revenue and number of employees has limited minimum impact on the predictive accuracy. This insight helps in framing the cyber insurance policy by placing more weight on the incident type.

3.4.3 Estimating the financial cost of cyber incidents

The final predictive model involves the financial cost incurred by a company following a cyber incident. It is reasonable to expect that this outcome variable is correlated with the number of affected entities. This turns out to be the case, the Kendall tau coefficient between the variables is 0.48 ($n = 10651$). Similar as before, the cost variable was discretized by splitting the values into three groups, representing no financial cost, cost between 0 – 100K USD and

costs larger than 100K USD (this value is somewhat arbitrary but gives a similar number of cases in the non-zero groups and as a result class imbalance is avoided). The accuracy of a Random Forests classifier in this case was 63.6% (SD 1.1%). Figure 5 (right panel) shows that the company sector is the most important feature for predicting the cost, while company revenue and number of employees has minimal importance, the latter of which was also observed when predicting the affected count in the previous section.

Figure 5: Feature importance scores for the number of affected entities (left) and total cost (right).

4 Conclusions

Malicious breach incidents dominate cyber threats across all sectors. Specifically, unauthorized contact/disclosure incidents mostly occur in finance, ASWR services and professional, scientific, and technical services. The services that appear to have the highest exposure are the ASWR and accommodation and food services. Phishing and malicious breach incidents involve the highest costs while cyber extortion incidents typically involve costs that are two orders of magnitude lower. Malicious breaches, most often target personal financial information. Personal identity information is mostly compromised in unauthorized contact/disclosure cases. These insights aid in designing cyber insurance policies that are sector and asset specific.

An analysis of the perpetrators of cyber incidents shows that hacktivists, foreign nation states and terrorists have similar targets and employ similar methods. In the same vein, employees, vendors, consultants and trusted third parties also exhibit similar adversarial tendencies. Though, the granularity of the classification of threat actors is coarse, this insight coupled with external information of potential adversarial interest in an organization's asset also aid in appropriately scoping the risk profile.

Though the Advisen cyber loss data feed is not sufficiently detailed to build a model that can predict the financial cost of a cyber incident with high accuracy, the model described here still provides useful insights into the factors affecting the financial cost of such incidents. The corporate sector in particular should be taken into account when determining insurance prices. A more accurate predictive model will require additional data. Combining the Advisen data set with other data sources is therefore of significant interest in the future. For example, data on corporate finance, stock price, news feeds and cyber threat indicators could be included.

References

Advisen. 2019. *Advisen Cyber Loss Data.* `https://www.advisenltd.com/data/cyber-loss-data/?utm_campaign=Cyber%20Data&utm_source=press-release&utm_medium=ireach`. [Online; accessed 15-September 2019].

Ashby, Simon, Trevor Buck, Stephanie Nöth-Zahn, and Thomas Peisl. 2018. "Emerging IT Risks: Insights from German Banking." *The Geneva Papers on Risk and Insurance - Issues and Practice* 43, no. 2 (April): 180–207.

Betterley, R. 2013. *Cyber/Privacy Insurance Market Survey 2013: Carriers Deepen Their Risk Management Services Benefits - Insureds Grow Increasingly Concerned with Coverage Limitations, online edition.* Last accessed 16 September 2019. `http://betterley.com/samples/cpims13_nt.pdf`.

Biener, Christian, Martin Eling, and Jan Hendrik Wirfs. 2015. "Insurability of Cyber Risk: An Empirical Analysis." *The Geneva Papers on Risk and Insurance - Issues and Practice* 40, no. 1 (January): 131–158.

Breiman, Leo. 2001. "Random Forests." *Mach. Learn.* 45, no. 1 (October): 5–32.

CISA. *Information Sharing Specifications for Cybersecurity.* `https://www.us-cert.gov/Information-Sharing-Specifications-Cybersecurity/`. Accessed: 2019-09-11.

Derrig, Richard, and Louise Francis. 2006. "Distinguishing the forest from the TREES: a comparison of tree based data mining methods." In *Casualty Actuarial Society Forum,* 1–49.

Dreyer, P., T. Jones, K. Klima, J. Oberholtzer, A. Strong, J. W. Welburn, and Z. Winkelman. 2018. "Estimating the Global Cost of Cyber Risk: Methodology and Examples." *RAND Corporation.*

Edwards, Benjamin, Steven Hofmeyr, and Stephanie Forrest. 2016. "Hype and heavy tails: A closer look at data breaches." *Journal of Cybersecurity* 2, no. 1 (December): 3–14.

Eling, M., and W Schnell. 2016. *Ten Key Questions on Cyber Risk and Cyber Risk Insurance.* Last accessed 16 September 2019. `https://www.genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public/cyber-risk-10_key_questions.pdf`.

Eling, M., and J. H. Wirfs. 2015. *Modelling and management of cyber risk.* Last accessed 16 September 2019. `https://www.actuaries.org/oslo2015/papers/IAALS-Wirfs&Eling.pdf`.

Eling, Martin, and Jan Wirfs. 2019a. "What are the actual costs of cyber risk events?" *European Journal of Operational Research* 272 (3): 1109–1119.

Eling, Martin, and Jan Wirfs. 2019b. "What are the actual costs of cyber risk events?" *European Journal of Operational Research* 272 (3): 1109–1119.

Guelman, Leo, Montserrat Guillen, and Ana M Perez-Marin. 2012. "Random forests for uplift modeling: an insurance customer retention case." In *International Conference on Modeling and Simulation in Engineering, Economics and Management,* 123–133. Springer.

Huth, M. 2018. "Economic impact of IoT cyber risk - analysing past and present to predict the future developments in IoT risk analysis and IoT cyber insurance." *IET Conference Proceedings* (January): 3–9.

Jacobs, V., J. Bulters, and M. van Wieren. 2016. "Modeling the Impact of Cyber Risk for Major Dutch Organizations." *Proceedings of the 15th European Conference on Cyber Warfare and Security* (July): 145–154.

Kesan, Carol M., Jay P. Hayes. 2017. "Strengthening Cybersecurity with Cyberinsurance Markets and Better Risk Assessment." *Minnesota Law Review* 102:191.

Kunreuther, Howard, and Mark V. Pauly. 2018. "Dynamic Insurance Decision-Making for Rare Events: The Role of Emotions." *The Geneva Papers on Risk and Insurance - Issues and Practice* 43, no. 2 (April): 335–355.

Lagazio, M., N. Sherif, and M. Cushman. 2014. "A multi-level approach to understanding the impact of cyber crime on the financial sector." *Computers & Security* 45:58–74.

Lu, X., D. Niyato, H. Jiang, P. Wang, and H. V. Poor. 2018. "Cyber Insurance for Heterogeneous Wireless Networks." *IEEE Communications Magazine* 56, no. 6 (June): 21–27.

Maillart, T., and D. Sornette. 2010. "Heavy-tailed distribution of cyber-risks." *The European Physical Journal B: Condensed Matter and Complex Systems* 75, no. 3 (June): 357–364.

NAICS. 2019. *NAICS Association.* https://www.naics.com/search-naics-codes-by-industry/. Accessed: 2019-09-11.

Orlando, Albina, Angelica Marotta, Stefano Nanni, Fabio Martinelli, and Artsiom Yautsiukhin. 2017. "Cyber - insurance survey." *Computer Science Review* (May): 35–61.

Pedregosa, F., et al. 2011. "Scikit-learn: Machine Learning in Python." *Journal of Machine Learning Research* 12:2825–2830.

Pooser, David M., Mark J. Browne, and Oleksandra Arkhangelska. 2018. "Growth in the Perception of Cyber Risk: Evidence from U.S. P&C Insurers." *The Geneva Papers on Risk and Insurance - Issues and Practice* 43, no. 2 (April): 208–223.

Refsdal, Atle, Bjørnar Solhaug, and Ketil Stølen. 2015. "Cyber-risk Management." In *Cyber-Risk Management,* 33–47. Springer International Publishing.

RMS. 2019. *Cyber Risk Landscape.* http://www.rms.com/models/cyber. [Online; accessed 15-September 2019].

Romanosky, Sasha. 2016. "Examining the costs and causes of cyber incidents." *Journal of Cybersecurity* 2, no. 2 (August): 121–135.

Romanosky, Sasha, Lillian Ablon, Andreas Kuehn, and Therese Jones. 2019. "Content analysis of cyber insurance policies: how do carriers price cyber risk?" Tyz002, *Journal of Cybersecurity* 5 (1).

Schnell, Werner. 2018. "Extreme Cyber Risks and the Non-diversification Trap." In *ARIA 2018 Annual Meeting.* August. https://www.alexandria.unisg.ch/255592/.

Shetty, Sachin, Michael McShane, Linfeng Zhang, Jay P. Kesan, Charles A. Kamhoua, Kevin Kwiat, and Laurent L. Njilla. 2018. "Reducing Informational Disadvantages to Improve Cyber Risk Management†." *The Geneva Papers on Risk and Insurance - Issues and Practice* 43, no. 2 (April): 224–238.

Smidt, Guido de, and Wouter Botzen. 2018. "Perceptions of Corporate Cyber Risks and Insurance Decision-Making." *The Geneva Papers on Risk and Insurance - Issues and Practice* 43, no. 2 (April): 239–274.

SwissRe. 2017. *Cyber: getting to grips with a complex risk.* `https://www.swissre.com/dam/jcr:995517ee27cd-4aae-b4b1-44fb862af25e/sigma1_2017_en.pdf`. [Online; accessed 15-September 2019].

Wheatley, Spencer, Thomas Maillart, and Didier Sornette. 2016. "The extreme risk of personal data breaches and the erosion of privacy." *The European Physical Journal B* 89:1–12.

Wu, Z., W. Lin, Z. Zhang, A. Wen, and L. Lin. 2017. "An Ensemble Random Forest Algorithm for Insurance Big Data Analysis." In *2017 IEEE International Conference on Computational Science and Engineering (CSE) and IEEE International Conference on Embedded and Ubiquitous Computing (EUC),* 1:531–536. July.