

# Towards Integrating Insurance Data into Information Security Investment Decision Making

Daniel W. Woods and Andrew C. Simpson  
Department of Computer Science, University of Oxford  
Wolfson Building, Parks Road, Oxford OX1 3QD  
United Kingdom  
firstname.lastname@cs.ox.ac.uk

**Abstract**—Making security investment decisions involves giving consideration to a variety of risks. However, there is little robust empirical evidence that can be used to support this process. This paper builds a road-map for incorporating cyber insurance data into existing security investment models. We propose an approach for using this data as an input for one investment model and introduce three distinct methods for evaluating the effectiveness of a new investment. We then describe a road-map for improving the insurance data collection process that aims to improve data utility for researchers. This approach could benefit those trying to justify an investment at all levels by providing evidence for the return on security.

**Keywords**—Risk assessment, data, empirical security, quantifying risk, quantifying impact, cyber insurance

## I. INTRODUCTION

The consequences of an organisation mismanaging security are increasingly severe. For example, a 2017 report [1] suggested that the average consolidated total cost of a data breach stands at \$3.62 million. Whatever one thinks of such calculations and conclusions, one thing is clear: security budgets are rising [2]. Some security researchers have asserted that cyber risks “cannot be managed better until they can be measured better” [3]. Yet, there has been little use of empirical data in this respect. Indeed, Verendel goes as far as suggesting that quantified security lacks validation by empirical results [4].

In this paper we propose a novel data source for evaluating security investments. We give consideration to how cyber insurance data might be applied to investment models that seek to quantify the benefits of different information security controls. We propose an approach for using this data within a pre-existing model. We also map out future steps to increase the utility of insurance data and move towards a scientific approach to investment models that incorporates insurance data, with a view to addressing the concerns aired by Verendel in [4].

The structure of the remainder of this paper is as follows. We start, in Section II, by providing the background for the rest of the paper, including a review of the related literature and an overview of how cyber insurance data is collected. Then, in Section III, we investigate how this data could be used in an existing investment model and introduces three approaches to estimating breach probability. In Section IV we discuss insurance data as an input, as well as the strengths of each of the three approaches. Section V looks at how insurance

processes can be standardised for the benefit of insurers and analysis alike. Finally, Section VI concludes the paper.

## II. BACKGROUND

In this section we introduce a popular investment model, before describing the format of insurance data and justifying why it provides a useful input for the model.

### A. Return on Security Investment

Investment models can be used to provide a cost-benefit analysis of a given investment. Return on Security Investment (ROSI) quantifies the benefits of a given investment adjusted for the cost of investment [5], which involves a calculation of the form

$$\text{ROSI} = \frac{\text{Benefits of Investment} - \text{Cost of Investment}}{\text{Cost of Investment}}$$

There have been numerous proposed methods of quantifying the benefit of investment. In [6], Böhme and Moore create a multiple round model that considers the benefits of delaying investment, which is extended to include penetration testing in [7]. In [8], the authors consider trade-offs between security and other factors (such as performance) in cloud computing. The need to quantify the impact of a successful cyber attack is common across these models.

In [9], Heitzenrater and Simpson estimate the benefits of security using the notion of Annualised Loss Expectancy (ALE), which is the product of the Annual Rate of Occurrence (ARO) and the Single Loss Expectancy (SLE). The Information Security Breaches Survey (ISBS) [10] — a self-reported survey of UK organisations — asks each respondent to report on their “worst-case loss”, which is used as an input for the SLE, with an annual rate of occurrence that is assumed to be 1. We adapt this approach to use insurance data, the format of which we consider in the next subsection.

### B. Insurance Data Format

Cyber insurance data has two important aspects for our approach: *ex-ante* assessment of the insured provides insights into which security controls are in place; *ex-post* claims financially quantify the effect of any cyber attacks suffered. The insurance industry’s ability to link the two across many insureds is a unique perspective.

TABLE I. THE RISKS AND COSTS COVERED BY EACH TYPE OF CYBER INSURANCE COVERAGE.

	Coverage	What It Covers
$A_1$	First-Party	The cost of replacing or restoring lost data. Excludes liability for intellectual property.
$A_2$	Data Privacy and Network Security Liability	Liability claims of a third party as a result of a data breach or an unintentional transmission of a computer virus.
$A_3$	Business Interruption	Revenues lost as a result of network down time.
$A_4$	Cyber-Extortion	Investigation costs, sometimes the extortion demand.
$A_5$	Public Relations	Fees for Public Relations firm to manage reputation in the event of a breach.
$A_6$	Multi-Media Liability	Costs relating to the content of a firm's website such as copyright infringement.
$A_7$	Professional Services	Liability relating to a service offer such as web hosting or internet service.

Claims may be differentiated according to the type of coverage offered. Table I details what is covered by a range of coverage types. Certain cyber risks are covered, while others are not. For example, the liability for a data breach is covered by  $A_2$ , while the risk of lost revenue as a result of systems downtime may be covered by  $A_3$  or  $A_4$ . However, the risk of lost revenue as a result of reputation damage is not generally covered. We do not claim this list to be exhaustive, not least because it will evolve over time.

Characterising what data is collected in the risk assessment is more difficult. Empirical research has identified conflicting results on how the risk assessment is conducted. Franke [11] reported that insurers choose to offer insurance based on the security assessment, Woods et al. [12] identified a range of security controls about which insurers collect information in the application process, while Romanosky et al. [13] showed that the majority of US insurers do not consider information security when pricing risk. This disparity makes it difficult to use the academic literature to identify the factors considered in risk assessment.

For this reason we will consider a proposed cyber insurance data standard [14] (claimed to be the first of its kind). In other lines of insurance, data collection has converged on a common standard<sup>1</sup>. The cyber insurance standard includes a collection of mandatory and optional categories relating to a policyholder — the mandatory categories mean that every policy inputted will contain, at a minimum, the firm's industry and revenue. We include a selection of the available fields in Table II.

In order to input scores such as the 'Privacy Policy Score' or the 'Encryption Quality Score', the data standard provides a Quality Score Rubric. This is a standardised evaluation of a number of aspects of an organisation's cyber security according to a pre-defined rationale. Scores range from a binary Yes/No choice through Low/Medium/High scoring to externally defined levels. For example the 'Anti Virus Quality Score' depends on the scope of installation, which is scored with 'Few workstations', 'Most workstations' and 'All endpoints' corresponding to poor, average and excellent respectively.

### C. Suitability of Insurance Data

We now attempt to motivate the use of insurance data within security investment models. In order to do so, we first look at one alternative — that of survey data. Of the surveys mentioned so far in this paper, we find that [1], [2] and [10] have 383, 260 and 1,125 responses respectively. Florêncio

and Herley [15] identify a number of issues with cyber-crime surveys; increasing the sample size is suggested as one solution to the problems associated with surveying rare phenomena such as data breaches. However, these numbers are low because (in part) surveys are expensive to conduct, whereas insurance data is created in proportion to the size of the market — as a by-product of normal business operations.

Conflicting cyber crime survey results is another issue that has been identified [16], which is (in part) explained by the different populations surveyed. This raises a question about the relevance of the losses of a large health care provider to a small retail firm. Meanwhile, industry code is mandatory in the data standard. In addition, there are optional fields including Founding Year, Employee Count and Revenue [14], providing the ability to focus only on data points 'relevant' to a given organisation. This would allow the "finer-grained breakdown of the data", in particular relating to employee numbers, called for in [9].

The difficulty of "mapping incidents to losses" is raised as a challenge in [17]. Utilising insurance data skips the incident step by solely considering direct losses. Claims financially quantify the effect of cyber attacks suffered — if a breach leads to no harm, then it is not reflected in claims data. Measurement is done in a unit of account that facilitates comparison. We acknowledge that fraudulent claims do weaken the data set.

As a firm's cyber security practices are generally considered sensitive information, surveys rarely collect data relating to this. However, this information is shared with insurers as they cannot accept liability for a risk without understanding the cyber security measures mitigating that risk. Importantly, the data standard of [14] provides a standardised format that allows for comparison between different insurer's assessments, which tend to be different to each other [13, 18].

Coverage being split into different policies delimits different losses. This is important because security controls do not affect these losses equally. For example, regularly backing-up a system mitigates the risk of a suffering a cyber-extortion attack, but will do little to protect against liability for a data breach. We propose a model to separate out these losses and consider how a given security control mitigates each loss specifically, which allows for a more fine-grained understanding of its effectiveness.

## III. INCORPORATING INSURANCE DATA INTO ROSI

The analysis of Heitzenrater and Simpson described in [9] defines the benefit of an investment in terms of the difference between  $ALE_s$  and  $ALE_0$ , the expected loss with and without

<sup>1</sup>See, for example, ACORD's data standards: <https://www.acord.org/standards/downloads/Pages/default.aspx>.

TABLE II. A SELECTION OF THE FIELDS OF THE AIR CYBER EXPOSURE DATA STANDARD, WITH MANDATORY FIELDS IN BOLD TEXT.

Section	Field
Organization	Breach History, Employee Count, Founding Year, <b>Industry Code</b> , Mergers/ Acquisitions, Ownership Type, <b>Revenue</b> , IT Maturity Score, BitSight Rating, ISO 27001 Indicator, Privacy Policy Score
Organization Data	<b>Data Type</b> , Record Count, Cost Per Unit, Data Back-up Frequency, Recovery Cost, Health Indicator
Asset	<b>Asset Type</b> , Business Interruption Cost, Recovery Cost, Location, Asset Count, Physical Security Measures, Anti-Virus Quality Score
Transfer	Business Interruption Cost, Recovery Cost, <b>Transfer Type</b> , Access Level, Payment Processor, DNS Provider, Encryption Quality Score, Cloud Type

a given security investment,  $s$ , respectively. This reduces to a binary choice between a set loss,  $\lambda$ , with probability,  $p$ :

$$ALE = p \cdot \lambda$$

Here,  $p = 1$  if no investment takes place because the set loss,  $\lambda$ , is equal to the average worst-case loss in the survey. If an investment is made, then  $p_s$  varies between 0 and 1 according to how effective the control is.

We propose how cyber insurance data can feed into  $\lambda$  and  $p$  in turn, and provide three different methods for estimating the probability of breach,  $p$ . We use a hypothetical example to illustrate this.

Consider a retail firm with around 50 employees looking at an investment with known cost,  $c$ . The investment is a phishing awareness scheme, which involves simulating a phishing attack against the firm's employees and giving further training to those who fall victim to it. Finally, suppose that there is a database X that contains insurance policies entered in a way that is consistent with the data standard of [14], complete with corresponding claims data delimited as in Table I. For simplicity, we will only consider coverage related to first-party costs, cyber extortion and multi-media liability ( $A_1$ ,  $A_4$  and  $A_6$  respectively).

#### A. Set Loss, $\lambda$

Claims data can be input as the set loss  $\lambda$ . This involves querying the database X to produce all results of entries with 'Industry Type = Retail' and '25 < Employee number < 100'. We then take an average of the total claims over a given time period to give  $\lambda_0$ . If the average amount claimed under first-party  $A_1$  is 12000, cyber extortion  $A_4$  is 16000 and multi-media liability  $A_6$  is 3200, then we will have the vector

$$\lambda_0 = (12000, 16000, 3200)$$

This is the expected loss if a breach occurs and with no further investment we assume that each will occur with probability 1. Thus

$$ALE_0 = \lambda_0 \cdot (1, 1, 1) = 31200$$

Next we look at calculating the probability of breach,  $p$ , if the phishing awareness scheme takes place.

#### B. Probability of Breach, $p$

The probability of breach,  $p$ , is a vector of the form

$$p = (p_1, p_2, p_3) \in [0, 1]^3$$

For example,  $p_1$  is the likelihood of suffering the set loss  $A_1$ . We introduce three methods of determining the value

of  $p$ : subjective effectiveness,  $p_{sub}$ , harnesses an individual's expectations; external effectiveness,  $p_{ext}$ , incorporates external data sources; and actuarial effectiveness,  $p_{act}$ , utilises data that relates losses to the controls that insureds have in place.

1) *Subjective Effectiveness*: This method relies upon an individual's judgement in classifying a control into different categories describing a percentage of risk mitigated. We outline reasoning behind how an anti-phishing awareness campaign mitigates each of  $A_1$ ,  $A_4$  and  $A_6$  to illustrate this method.

Making an arbitrary choice, suppose that High, Low and Ineffective correspond to mitigating 75%, 25% and 0% of the risk respectively.

First-Party Coverage covers the costs of restoring or replacing lost data. Many of these losses are caused by internal errors, which a phishing campaign does not mitigate. Consequently, we say that the control has a low effect on mitigating the risks covered by First-Party Coverage. Conversely, the majority of cyber extortion attacks originate from outside the organisation. Consequently, we conclude that the control will have a high effect on  $A_4$ . Finally, phishing awareness will have a negligible effect on  $A_6$ , which covers liability for content on the organisation's media outlets.

Given the above, we have the following.

$$\begin{aligned} p_{sub} &= (0.25, 0.75, 0) \\ ALE_{sub} &= (0.25, 0.75, 0) \cdot (12000, 16000, 3200) \\ &= 15000 \\ ROSI_{sub} &= \frac{ALE_0 - ALE_{sub} - c}{c} \end{aligned}$$

2) *External Effectiveness*: The 'external effectiveness',  $p_{ext}$ , uses data provided by external sources to quantify the effectiveness of an investment in mitigating the expected loss under a given coverage,  $A_k$ . Effectiveness ( $\alpha_{k_i}$ ) is likely measured against a particular attack,  $T_i$ . This must be corrected for the proportion of losses ( $\beta_{k_i}$ ) covered by  $A_k$  that can be attributed to the attack,  $T_i$ . We must then sum over all the different attacks that lead to a given loss, weighted by the proportion of attacks faced. This results in the following.

$$\begin{aligned} p_k &= 1 - \sum_i \alpha_{k_i} \beta_{k_i} \\ p_{ext} &= (p_1, p_2, \dots, p_7) \end{aligned}$$

Ideally, we would delimit all possible attacks so that  $\sum_i \beta_{k_i} = 1$ .

Suppose that we failed to list some attack,  $T_i$ . Then, the control's effectiveness  $\alpha_t$  in mitigating that attack would not be reflected in the sum  $p_k$ . The result would be an under-estimate of the effectiveness of the control.

To illustrate this method, we make the assumption that if  $x\%$  of simulated phishing attacks succeed, then  $x\%$  of real attacks will succeed. With  $\alpha_{k_i}$  signifying the change before and after an awareness campaign, we suggest the following as appropriate inputs:

$$\begin{aligned}\alpha_{k_i} &= \Delta \text{Percentage of Successful Simulated Attacks} \\ \beta_{k_i} &= \text{Proportion of } A_i \text{ losses involving phishing}\end{aligned}$$

One source of  $\beta_{k_i}$  is reports of the percentage of attacks involving phishing, with such a report from 2012 suggesting that 90% of targeted attacks start with phishing [19]. Now suppose, with  $\lambda_0 = (c_1, c_2, c_3)$ , that phishing attacks  $T_1$  result in 90% of the losses  $c_1$  and  $c_2$ , but none of the multi-media losses  $c_3$ . If the hypothetical phishing awareness campaign led to a fall from 80% to 40% of phishing attacks, then we would have

$$\begin{aligned}\alpha_{1_1} &= \alpha_{2_1} = \alpha_{3_1} = 0.4 \\ \beta_{1_1} &= \beta_{2_1} = 0.9 \text{ and } \beta_{3_1} = 0\end{aligned}$$

with  $\alpha_{k_i} = 0$  for  $a \neq 1$  as the phishing awareness campaign only affects phishing attacks. These result in an effective reduction of 0.36 across each  $A_1$  and  $A_3$  and no reduction of  $A_6$  so that

$$\begin{aligned}p_{ext} &= (0.64, 0.64, 1) \\ ALE_{ext} &= (0.64, 0.64, 1) \cdot (12000, 16000, 3200) \\ &= 21120 \\ ROSI_{act} &= \frac{ALE_0 - ALE_{ext} - c}{c}\end{aligned}$$

3) *Actuarial Effectiveness*: The ‘actuarial effectiveness’ is named as such due to the similarity to approaches found in actuarial science. These involve delimiting a series of insured individuals into categories, then looking at the expected result across each category. For example, a mortality table represents the survivorship of people from each age range. As they age, people ‘move’ between age categories and their life expectancy changes accordingly.

In a similar way, for a security control,  $C$ , we can divide the sample into two categories: those who do not have the control and those who have the control, with corresponding average losses  $\lambda_0$  and  $\lambda_1$  respectively. A control here may consist of an individual control or a collection thereof — recognising that a collection may capture the interdependence of different controls, but will likely result in small sample of organisations with that collection.

We assume that by implementing the control  $C$  an organisation’s new expected loss is the average of all organisations who have that control. This assumption leads to

$$\text{Benefits of Investment} = \lambda_0 - \lambda_1$$

With  $\lambda_0 = (c_1, c_2, c_3)$  and  $\lambda_1 = (c'_1, c'_2, c'_3)$ , we set  $p_{act}$  as

$$p_{act} = \left( \frac{c'_1}{c_1}, \frac{c'_2}{c_2}, \frac{c'_3}{c_3} \right)$$

to ensure that this result holds.

Note that  $ALE_{act} = \lambda_1$ , so that

$$ROSI_{act} = \frac{\lambda_0 - \lambda_{act} - c}{c}$$

## IV. DISCUSSION

In this section we discuss the previous example, giving consideration to the strengths and weaknesses of each approach.

When calculating  $\lambda$ , there is a trade-off between the similarity of firms in the sample and the size of the sample. A large sample harnesses the law of large numbers, while a smaller sample guarantees that the other data points in the sample are relevant. For example, historical data becomes less relevant to predicting future claims with time.

The ‘subjective effectiveness’ approach carries the stigma associated with subjectivity. However, the skeptical reader should remain cognisant that risk management often splits the probability of breach into low, medium and high probability based on individual judgment [20]. The security professionals carrying out the risk assessment possess knowledge and experience of the particular organisation that may be relevant.

Providers of the data used in the ‘external effectiveness’ approach tend to aggregate data from a wide range of organisations and contexts, arguably losing relevance to the particular circumstances. The effectiveness of a control  $\alpha$  is often produced by firms in the security industry, without peer review and (arguably) with an agenda. Meanwhile, academic studies tend to involve experiments, which do not adequately simulate ‘the wild’. Estimating the proportion of losses  $\beta$  that can be attributed to a particular attack is non-trivial. In addition, the challenge of listing all of the relevant attacks a given control mitigates. Further, even if a control effectively mitigates an attack, it is not clear that another attack vector would not instead be used by the attacker.

The ‘actuarial effectiveness’ approach assumes that, by implementing a control  $C$ , an organisation’s new expected loss is the average of all organisations who have that control, which may attribute causation to correlation. Considering individual controls obscures how controls relate to each other, which was identified as a challenge in cyber insurance assessment in [18]; the marginal benefit of a second firewall is smaller than putting in place the original, for example. Selecting sets of controls results in small sample sizes that are relatively more influenced by the random nature of cyber losses. Further, only considering existing data may discentivise novel controls in which the sample size of organisations with the control is too small to gain a meaningful expected loss. This could lead to herding behaviour as organisations converge on the practices of the best risk class. This would, of course, stifle innovation.

These strength and weaknesses must be considered in the context of the security control being evaluated. Consequently, the method of calculating breach probability can be chosen according to the investment and, importantly, according to the decision maker. For example, we believe that the phishing awareness campaign suits the ‘external approach’ because it accurately simulates ‘the wild’. Another decision maker may possess experience with the control in question and so the ‘subjective approach’ might be more appropriate. Finally, the data might be a particularly good fit for a particular investment, so that the ‘actuarial approach’ is appropriate. Our approach allows for a ‘mix and match’ approach, in which an investment quantified using a subjective assessment can be compared with one that is grounded in data.

## V. INDUSTRY ROAD-MAP

In this section we outline a brief research agenda to maximise the utility of cyber insurance data for decision making. Much of this relates to data collection processes, as standardisation enables collaborative research and large-scale analytics.

1) *Standardisation*: Meland et al. [21] identified a “lack of standardized indicators” when assessing cyber risk. The ex-ante assessment process involves applicants filling out self-assessed forms that vary greatly between organisations [13, 18]. If these forms followed a common standard, then they could provide risk profiles for the actuarial effectiveness approach: organisations could be grouped according to the results of the assessment process.

Another area that would benefit from standardisation is the wording of the coverage, which defines the risks that the policy covers. Standardising the wording would give a clearer mapping between the amount claimed and the losses faced; this delineation of types of loss is one of the strengths of insurance data. If coverage does not extend to risks not currently covered, like intellectual property and reputation damage [22], then the approach we outline cannot account for these risks.

2) *Ex-Post Forensics*: The standardised collection of forensic evidence could allow researchers to link threats and vulnerabilities to losses. At present, this evidence is largely collected by security vendors who often will not have a complete picture regarding the impact of a given security event. One solution is to build forensics into the claims process, either through loss adjusters or via a forensics section in the claims form. This data could then be used as an input for the external approach.

For example, if  $x\%$  of data breaches involved human engineering, we could better assess investing in employee training — especially if the data revealed that one area of coverage was more vulnerable to a given attack than another. This forensic evidence could also help settle exclusions disputes and provide a mechanism to combat moral hazard, which is a problem identified in [23].

3) *Policy Landscape*: Many initiatives that support our road-map are already being discussed by various stakeholders. Standardisation was identified as a policy objective by both policy-makers and industry bodies [12]. For example, Lloyd’s of London has led the development of core data requirements that both “AIR and the RMS/Cambridge team” agree upon [24]. In addition ACORD, a global insurance standards body, has begun a project to standardise proposal forms [25].

In addition, there is political will to make insurance data available for academic research. This can be seen in the 2015 report by the UK Government stating that “The Government will work together with the insurance industry, including the ABI and Lloyd’s, to establish a forum for data and insight exchange and for policy discussions” [26]. These initiatives highlight that there is pre-existing buy-in from stakeholders.

## VI. CONCLUSION

We have constructed a hypothetical example to illustrate how cyber insurance data can be used to produce ROSI calculations. This involves using claims data as an input for the

‘set loss’ parameter and classifying three distinct approaches to an input for ‘probability of breach’. We also described how future steps might improve insurance processes to increase data utility, linking this to a number of current policy initiatives.

The discussion identified a number of weaknesses in each of the approaches to estimating the effectiveness of a given control. Future work could identify alternative data sources for *external effectiveness* and develop a clearer heuristic for *subjective effectiveness*. Real insurance data may determine whether *actuarial effectiveness* is a viable option.

The road-map provides an opportunity for further research; developing a standardised forensics section that will collect data in a way that is useful for the external effectiveness approach and for wider academic research is an important next step in this respect.

## ACKNOWLEDGMENT

The authors thank the anonymous reviewers for their helpful and constructive comments. Daniel Woods’ research is funded by the EPSRC via the Centre for Doctoral Training in Cyber Security at the University of Oxford.

## REFERENCES

- [1] Ponemon Institute, “Cost of Data Breach Study Available: <http://www-03.ibm.com/security/data-breach/>,” 2016, [Online; accessed 27-June-2016].
- [2] The IISP, “Security market trends and predictions,” 2015, [Online; accessed 27-June-2016]. [Online]. Available: <http://iisp.informz.net/IISP/WhitePaper.pdf>
- [3] R. Anderson and T. Moore, “The economics of information security,” *Science*, vol. 314, no. 5799, pp. 610–613, 2006.
- [4] V. Verendel, “Quantified security is a weak hypothesis: A critical survey of results and assumptions,” in *Proceedings of the Workshop on New Security Paradigms Workshop*. ACM, 2009, pp. 37–50.
- [5] W. Sonnenreich, J. Albanese, and B. Stout, “Return on security investment (ROSI)-a practical quantitative model,” *Journal of Research and Practice in Information Technology*, vol. 38, no. 1, pp. 45–56, 2006.
- [6] R. Böhme and T. Moore, “The iterated weakest link,” *IEEE Security & Privacy*, vol. 8, no. 1, pp. 53–55, 2010.
- [7] R. Böhme and M. Félegyházi, “Optimal information security investment with penetration testing,” in *International Conference on Decision and Game Theory for Security*. Springer, 2010, pp. 21–37.
- [8] N. Tsalis, M. Theoharidou, and D. Gritzalis, “Return on security investment for cloud platforms,” in *Cloud Computing Technology and Science (CloudCom)*, 2013 IEEE 5th International Conference on, vol. 2. IEEE, 2013, pp. 132–137.
- [9] C. D. Heitzenrater and A. C. Simpson, “Policy, statistics, and questions: Reflections on UK cyber security disclosures,” *Journal of Cybersecurity*, vol. 2, no. 1, pp. 43–56, 2016.
- [10] Department for Business, Innovation & Skills, “Information security breaches survey,” 2015, [Online; accessed 27-July-2016]. [Online]. Available: <https://www.gov.uk/government/publications/information-security-breaches-survey-2015>
- [11] U. Franke, “The cyber insurance market in Sweden,” *Computers & Security*, vol. 68, pp. 130–144, 2017.
- [12] D. Woods and A. C. Simpson, “Policy measures and cyber insurance: A framework,” *Journal of Cyber Policy*, vol. 2, no. 2, pp. 209–226, 2017.
- [13] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones, “Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk?” in *Proceedings of The 16th Workshop on the Economics of Information Security (WEIS 2017)*, 2017.
- [14] V. Analytics, “Verisk Cyber Exposure Data Standard Available: <http://www.air-worldwide.com/Documentation/Cyber-Exposure-Data-Standard/Index.htm>,” 2016, [Online; accessed 27-June-2016].

- [15] D. Florêncio and C. Herley, "Sex, lies and cyber-crime surveys," in *Economics of information security and privacy III*. Springer, 2013, pp. 35–53.
- [16] S. L. Pfleeger and R. Rue, "Cybersecurity economic issues: Clearing the path to good practice," *IEEE Software*, vol. 25, no. 1, pp. 35–42, 2008.
- [17] R. Böhme and G. Kataria, "Models and measures for correlation in cyber-insurance," in *Proceedings of Workshop of Economic Information Security (WEIS)*, 2006.
- [18] D. Woods, I. Agraftiotis, J. R. Nurse, and S. Creese, "Mapping the coverage of security controls in cyber insurance proposal forms," *Journal of Internet Services and Applications*, vol. 8, no. 1, p. 8, 2017.
- [19] Trend Micro (UK) Limited, "Over 90 Percent of Targeted Attacks start with Spear Phishing Emails. Available: <http://www.trendmicro.co.uk/newsroom/pr/over-percent-of-targeted-attacks-start-with-spear-phishing-emails/>," 2012, [Online; accessed 27-June-2016].
- [20] T. R. Peltier, *Information security risk analysis*. CRC press, 2005.
- [21] P. H. Meland, I. A. Tøndel, and B. Solhaug, "Mitigating risk with cyberinsurance," *IEEE Security & Privacy*, vol. 13, no. 6, pp. 38–43, 2015.
- [22] C. A. Siegel, T. R. Sagalow, and P. Serritella, "Cyber-risk management: technical and insurance controls for enterprise-level security," *Information Systems Security*, vol. 11, no. 4, pp. 33–49, 2002.
- [23] N. Shetty, G. Schwartz, M. Felegyhazi, and J. Walrand, "Competitive cyber-insurance and internet security," in *Proceedings of Workshop of Economic Information Security (WEIS)*. Springer, 2010, pp. 229–247.
- [24] Lloyd's of London, "Cyber Core Data Requirements Available: <http://www.lloyds.com/news-and-insight/risk-insight/emerging-risks-team/cyber-core-data-requirements/>," 2016, [Online; accessed 27-June-2016].
- [25] ACORD, "Evolving the ACORD Standard to include Cyber Liability," 2015, [Online; accessed 27-July-2016]. [Online]. Available: [\url{https://www.agencyport.com/blog/evolving-the-acord-standard-to-include-cyber-liability/}](https://www.agencyport.com/blog/evolving-the-acord-standard-to-include-cyber-liability/)
- [26] HM Government & Marsh Ltd, "UK Cyber Security: The Role of Insurance In Managing and Mitigating the Risk," 2015.