British Actuarial Journal (2019), Vol. **24**, e20, pp. 1–15
doi:10.1017/S1357321719000138

ABSTRACT OF THE LONDON DISCUSSION

Cyber operational risk scenarios for insurance companies

Abstract of the London Discussion

[Institute and Faculty of Actuaries, Sessional Research Event, London, 15 October 2018]

This abstract refers to the following paper: Egan, R., Cartagena, S., Mohamed, R., Gosrani, V., Grewal, J., Acharyya, M., Dee, A., Bajaj, R., Jaeger, V.-J., Katz, D., Meghen, P., Silley, M., Nasser-Probert, S., Pikinska, J., Rubin, R. & Ang, K. (2019). Cyber operational risk scenarios for insurance companies. *British Actuarial Journal,* 24: e6. doi:10.1017/S1357321718000284

The Chair (Mrs L. Williamson, F.I.A.): Today's sessional research xevent is Cyber Operational Risk Scenarios for Insurance Companies. I am a director at Willis Towers Watson, in a current role of managing the risk and compliance for LifeSight Master Trust. I am a member of the Risk Management Board at the IFoA. An important part of what we do is championing people who work in risk-related and wider fields and who look to the IFoA for support and resource in risk-related areas.

We have three sub-committees, one of which is the Research and Thought Leadership Committee, who work directly with our working parties. We have 6–8 working parties at the moment, one of which is presenting their findings today.

While I work in the pensions field, one of the things that I have noticed is how important cyber has been in my day-to-day job over the past few years. When I look back over 5 years ago, the biggest concern in relation to cyber was someone hacking your Hotmail account and sending strange emails, but now the stakes have changed and are much higher.

As the technology is evolving so rapidly, it is challenging to see all the threats that are out there and try to pull together a strategy to combat them. Today's paper goes some way towards helping address this. It is doing what we, as actuaries, like to do. It is finding a way where we can identify problems and a way to assign risk, a way that we can look at probabilities and redraw and quantify the risks that we are facing. This fits in quite nicely with the challenges that have gone before it when you look at operational risk more widely.

We have three speakers here. The first speaker is Rory Egan, who is chair of the cyber risk investigations working party. He is a senior cyber actuary at Munich Re and part of Munich Re's corporate underwriting division focusing on accumulation risk management, data and pricing for cyber risk in insurance and reinsurance contracts.

Our second speaker is Ramiz Mohamed, who is head of cyber pricing at Hiscox. He is part of their central cyber business unit with focus on creating coordination, consistency and capability and pricing and data across the group.

Finally, we have Vanessa Jaeger, who is a principal consultant at Aon. She is a scheme actuary and leads a retirement cyber risk consulting specialist team providing guidance to the trustees and sponsors of DB and DC pension schemes in understanding and mitigating cyber risk and developing instant response plans.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

Mr R. Egan, F.I.A.A., C.E.R.A. (introducing the paper): The purpose of the working party's research is to provide insight into actuaries working on capital requirements for insurers, setting out the potential impact of cyber risk events and the measures available to mitigate this risk. The aim is to create a greater awareness of the risks for insurers, and highlight emerging issues in an area that is changing rapidly as the dependency on computer systems to support insurers' business increases.

We have 16 volunteers on the working party. We are not all actuaries but most of us are. We currently have one cyber risk expert, who has been invaluable to us as we developed this paper.

I am going to set the scene and explain why this is an important topic for us in the greater context of the cyber risk landscape, and then Ramiz (Mohamed) and Vanessa (Jaeger) will go into the detail about how we approached it in the paper in terms of a framework for looking at operational risk from cyber. We will then cover some of our worked examples of scenarios and the process of quantifying potential costs from these incidents.

I will then close, leading to discussion and Q&As. We will be interested to hear what you think we should do next, what further developments you would like to see on this topic.

First, starting with our definition of cyber, for the working party, we are currently keeping the definition extremely broad. It is described as the risk of doing business in the digital age. That includes the obvious like cyber-attacks, hacking and denial of service attacks, and also extends to more general IT operational failures to include fat-finger syndrome which is attributable to human error, such as pressing the wrong button leading to data being released or systems failing. We have a broad term for cyber that we work within which includes both malicious cyber events and non-malicious incidents with similar outcomes. We focus on operational risk.

Ramiz and I work in the cyber insurance market and (within the working party) we have a range of volunteers with different roles. We are finishing this phase of work around operational risk scenarios and want to hear from you what we could possibly tackle next. It could be taking this further or considering accumulation risk issues, for example, or a bit of both.

Cyber events are regularly reported in the news. We have picked out a few recent headlines; incidents that we looked at and thought "Could that happen at an insurance company?" And, could it be worse? We also wanted to tackle extreme scenarios, like the one in 200-year return period. Were these one in 200 events, or do we need to think of something even more extreme?

What I draw from some of these is the importance of how the company which is attacked, or suffers an incident, addresses the attack, and how much of a bearing that has on how they are affected at all stages, from planning before an incident occurs, what you do when an incident is discovered, in the immediate aftermath and throughout ensuing investigations. Those are all crucially important.

For example, the NHS could have prevented the ransomware attack (called "WannaCry") that they suffered. That attack was aimed at a particular version of Microsoft Windows operating systems. But a patch was available for several months before the attack, which the NHS could have applied and then they would never have suffered that incident in the first place. That is an example of the importance of risk mitigation.

I have heard from a colleague that the cost of the incident to the NHS has been estimated at £90 million in business disruption and rebuilding IT infrastructure.

To take another example, think of Uber. It was not unique in the fact that it suffered a data breach where customer data and also data on their employees, or drivers, was breached. What was perhaps unique in a negative way in their case was that they decided to conceal the fact that the data breach happened, which flies in the face of the regulatory requirements that they operate under. Now they are looking at about a $160 million fine for the decision to conceal.

Everybody is at risk from these attacks. How well you mitigate and react when it happens has a huge bearing on what happens to your bottom line.

If you go back to 2013–2014, the company attacks you heard about were targeted at individual companies including the company Target. Hackers were trying to steal the data, which they believed was interesting, or valuable, and from which they could make money in most cases. The common thread was that data was the asset of interest to the attackers, and that only one company was being affected, albeit in a massive way.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

Moving forward to 2016 and 2017, the Dyn attack, the WannaCry attack that affected the NHS, the NotPetya attack were of special interest to reinsurers, because these showed the catastrophic, or accumulation, potential that exists for cyber-attacks where the company that you work for, an insurance company perhaps, might not be in the cross-hairs of the hacker, but you just get caught up in some wider global attack. Also, the impact of those attacks was not about individuals having their data breached and needing to be compensated, but it was more about disruption or business interruption impact on those companies.

The cyber insurance market has had to respond in terms of the products that they are offering, shifting away from traditional data breach focused cyber insurance products to doing business interruption type of risk.

If anybody thought that there was a sea change in terms of the type of risk because of those accumulation events, the latest ones, Equifax and BritishAirways, show that the data breach on one single company is still happening.

Who is carrying out these attacks? And, more importantly, why do we care or should we care? It is a spectrum from the stereotypical kid in his parents' basement having a bit of a laugh with his mates: "Look, I have hacked into this company. Isn't that cool?" all the way up to nation-state sponsored groups of hackers with highly sophisticated cyber so-called "weapons".

If you are working for an insurance company, you might consider: "Do I need to worry about nation-states if I am doing business in just this region of the UK?" The answer is again "yes" because of the collateral damage aspect. You need to have that at the top of the list to consider when looking at what scenarios could impact your business. Those weapons are just so powerful.

On the other end of the spectrum: "Do I need to worry about a kid in his parents' basement?" The tools that he or she is using are also increasing in power, so they can make a ransomware attack against your company without really knowing how to code. They can hire ransomware tools at low cost and ransomware your company.

You might rather be dealing with serious organised crime people on the other side rather than kids because at least they know how to do it properly. If you decide to pay a ransom, then maybe you will get your data back; but maybe the kid does not know what he is doing and you pay him the ransom and then he does not have the skills to decrypt your data and send it back to you. Whether you pay the ransom or not could be the least of your concerns if you are dealing with an incompetent hacker actor on the other side. We have also considered the threat of insiders and will talk about these in a bit more detail.

How are they carrying out attacks? It varies and I take from this that we need to find a way to defend against these attacks. There is a range of different types of defences that we look at in the scenarios and in the framework that we recommend. You do need cyber security expertise and tools, the standard things like firewalls, patching, and all that kind of security technology. Often it is the human that is being hacked not the computer, so you need to tell your employees what phishing is, and how to avoid it. You also need to stop the worst employee, that is the one that is going to click on every email and every link, in your company so that you can improve that weakest link.

The combination of human factors, technology factors and the risk mitigation framework that you need to put together will need to encompass that whole spectrum of risks.

How big of a risk is cyber within the broader operational risk category or landscape? Is this overblown? Should we even exist as a working party? Allianz carried out a global survey. It is called the global risk barometer: they interviewed risk managers from all industries in all countries and they asked: "What is keeping you up at night? What are you worried about that could affect your business; that could take your business down; that could cost you money?"

What is in scope for us is insurance companies in the UK. Insurance comes within the category of financial services. Those risk managers identify cybercrime as the number-one type of risk that they are worried about. For companies in the UK, across all industries, the same story, cyber incidents are the number-one topic within the category of operational risk that they are concerned about.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

As actuaries we have a role to play in terms of the quantification piece of the risk management approach, by determining the capital that our companies need to hold against cyber risk. Additionally, there is going to be a lot of regulatory interest in cyber, given its importance. We already see that with new legislation such as General Data Protection Regulation (GDPR), and the Prudential Regulation Authority (PRA) and the Pensions Regulator, others are taking an interest in that. If they are not knocking on your door already, then you can expect that soon.

What we are putting forward here is a framework in terms of how to quantify the risk. It is not one size fits all: we have worked examples of scenarios, but you will need to tailor the approach to your company. The framework is arguably more important than the individual scenarios that we have quantified. We have gone through that process and learnt a lot from that in terms of the challenges that there are.

Miss V. J. Jaeger, F.I.A. (introducing the paper): How did we go about our research? First, the group had a brainstorming session with the aim of identifying a range of operational risk scenarios across different types of insurance products, considering areas such as IT systems, volumes and types of data stored.

This led us to consider a number of scenarios. We looked at a range of root causes, with the common theme of data extortion and system compromise. We wanted to look at scenarios impacting both the general and life insurers as well as emerging technology.

We reduced this to seven. Then we all took a vote to decide on the final three. This outcome led us looking at three scenarios which were: employee-linked data at a general insurer; cyber extortion at a life insurer; and a motor insurer telematics device hack. Once we had all these three together, we then fleshed out the details to find a tangible narrative to take this forward and then identify the cost.

As Rory Egan indicated, we wanted to consider the one in 200 event, so extreme but still plausible.

We then decided to break into three groups to investigate each scenario. Although we had agreed on the framework to present the output, each group took the task differently, which led to different ways of looking at the problem. This had its benefits as it included wide discussions about the output and meant that we needed to review the output for consistency, making further adjustments based on our discussions.

This clearly highlighted that there is no set approach when it comes to assessing operational risk. We all agree that, on reflection, if we were doing this research again, we would look to set a consistent approach first so as to remove some of the later discussions.

But there is that challenge with cyber threat in that they are always evolving and there is not always consistency in things like the terminology used and the approach taken for assessing this risk.

We had considered whether to try to establish a new taxonomy for this assessment. We concluded that there are a number of different frameworks out there already. These are perfectly suitable for our needs. For anyone undergoing this type of assessment, there is certainly value in selecting a specific framework to use early on and then sticking to this.

What was this framework? We selected to use the categorisation method for cyber risk which was proposed in the CRO Forum concept paper. The aim of this paper is to assist with the data capture for cyber incidents. We use this to breakdown the cause and the impact for each scenario.

The other taxonomy we used was the voluntary National Institute of Standards and Technology (NIST) cyber security framework. This framework has been developed to provide standards, guidelines and best practice to manage cyber security–related risks.

While it is a guide for US private sector organisations, it is used by organisations across the world. It is one that most gravitates when we are discussing cyber risks. The NIST framework comprises 22 control categories across five core functions to ensure that a company responds to cyber risk. These core functions are identify, protect, detect, respond and recover.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

We used this to assess the potential impact of the event and failure of a control, considering both the frequency of the event and the severity. We used version 1.0 for our assessment. We realised that when we were going through this, there was a new version as well, version 1.1. We did do that check to make sure that there was no material impact. This shows that everything is still evolving, even while we are doing our research.

We do recognise that there are other frameworks out there such as Cyber Essentials ISO 27001, and these may be equally suitable for your assessments.

Once that common taxonomy was agreed, we looked at the tangible and intangible losses which could arise from the failure in the cyber-related processes under each scenario. Here we explored key issues such as business impact on high-value assets and key weaknesses or dependencies.

Next, we considered what the important aspects are for building a framework. We concluded that the most important aspects are consistency, in particular being able to repeat any assessment, so following new data, new experiences and better knowledge. Cyber risk will continue to evolve, as will the taxonomy, the views on assumptions and publicly available data and the cost of cyber attacks, so it is important to develop an approach which can be repeated for different scenarios, but also one that you can continue to update.

The consistency also helped the communications as it is easier to communicate over time if the terminology is remaining consistent.

We then considered in more detail threat actors and vectors and whether these should be assessed. This is important as we need to understand the operating environment and know the specific threat. So whether the actors are insiders, hacktivists or organised criminals and the form of the attack, so the threat vector. This helps a company to consider how an attack may occur.

The next area is about quantification. We wanted to ensure consistency across all scenarios, for example, with data recovery cost, and an approach which was transparent, recognising the need to be able to justify and articulate the quantification process.

We have also set out our rationale for costings.

The final stage is the validation process. For our research, each team reviewed each other's scenarios output, and we also provided challenge for these costings based on our own experience. There were significant discussions on the sources of information available to validate the costings, and the range and limitations of available data sources, especially when you are looking at breaking down the cost of a cyber-attack in specific elements. We sought validation from cyber experts. We had lots of discussion over whether to include costs as a specific figure or just to consider the percentage of the overall costs. It took several discussions before we settled on our final output.

The fact that a group of 16 had a lot of discussion to reach a conclusion demonstrates the uncertainty with reaching a reasonable solution.

The purpose of the paper is to support actuaries with developing their own scenarios and accessing their own circumstances. Our scenarios are purposely designed to support your thinking. We are aware that you may hold different views and identify different outcomes.

Mr R. Mohamed, F.I.A. (introducing the paper): Setting out our scenarios within the CRO Forum framework helped us compare events between the different groups. The CRO Forum proposed a methodology for categorising cyber risk. The proposal was originally designed to assist with data capture on cyber claims, but we found it a useful framework to base our scenarios on.

The paper sets out four steps in a cyber event. First the root cause, that is, the underlying motive or the reason for the event. Second, the threat and who did it and how they did it. Third, the incident itself, and what happened to the company. Finally, the impact of the cyber attack and how much did it cost?

Each step breaks down into further categories which allowed us to build a full picture for each of our scenarios. For example, the root causes breakdown into people, processes and systems, and external factors. Within people, this breaks down into employee conduct, culture and behaviour,

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

and employee deliberate harmful acts, to name but a few. We were certainly keen to include one of the scenarios which was relating to internal employees. Cyber is often seen as an external threat. However, Verizon reported that last year 30% of cyber incidents arose from internal employees across all industries.

Within threat, we had a list of possibilities that could have carried out the attack. We have mentioned insiders. They can be those who are disgruntled and looking to cause embarrassment for the organisation or maybe just inflict damage on the system.

There are those who have criminal intentions, or it may simply be unintentional. Insiders could also extend to third parties such as consultants, or suppliers. The supply chain risk is increasingly seen as a risk that needs to be controlled and mitigated. There is also the Hollywood-style attacks, the hacktivists, the hackers, the organised crimes, and the most sophisticated, the nation-state attackers.

There has been a range of different actors and threat vectors in the insurance sector. Phishing emails from hackers are commonplace. According to PhishMe, even a well-trained, vigilant workforce can have a susceptibility rate of 5%. Social engineering is an increasingly popular way of getting through to companies. This is where the attacker uses psychological manipulation of people as a way of getting information or access to data or money. The poster child for this is an email from the CFO or from some purported vendor or asking for a payment to be wired to a specific bank account, often their bank account.

We have also seen nation-state involvement in the insurance sector as well. Rory Egan mentioned the non-targeted version. However, the Californian Insurance Commission has stated that it is likely the Anthem medical data breach in 2015 was caused by a foreign government.

Nation-state attacks are now a reality for insurers. The data they hold is valuable, particularly if the attacker believes that the data may relate to national infrastructure, key businesses, high net worth individuals or high-ranking government officials.

The next step is the incident. The incident can impact the confidentiality, integrity and availability of data and systems known as the CIA triad, as well as a system malfunction or misuse in criminal activity.

The final step is the impact step. We will look at it split down into a number of cross-components. We illustrate these through a discussion of our scenarios.

We intentionally chose very different scenarios to highlight the range of events insurers are vulnerable to. One common theme of these scenarios was it involved a malicious, targeted element. It is not only those malicious targeted attacks that can cause the highest cost. NotPetya last year and the British Airways IT failure, which grounded flights for 3 days, were examples of extremely costly events that were not targeted.

We have chosen an extreme but plausible scenario and looked at setting the costs of a one in 200 return period. However, we recommend starting with scenarios and costs at lower return periods and then extrapolating to a one in 200.

The overall cost that you will see to the hypothetical insurers ended up being very different. However, we want to stress the importance of the framework in this paper and the results demonstrate how the process could be applied. Our estimates include a significant degree of subjectivity, given the scarcity of data and the continually evolving threat and regulatory landscape. It is important to apply any scenarios to the specifics of your entities and document the rationale and supporting evidence for selection in order to track changes over time.

Scenario 1 is a disgruntled employee leaking policyholder data at a general insurer.

Data breach is one of the traditional events associated with cyber risk and for us in the insurance sector a very major risk.

In this scenario, we have assumed that the data was leaked from an internal member of staff who may have greater access to more sensitive data.

A real-life example of an internal data leak is Morrisons in 2014. A disgruntled employee in the internal audit team published names, addresses, phone numbers and pay details for thousands of

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

current and former employees. Last year, the UK High Court held Morrisons vicariously liable, which could lead to the first ever data leak class action in the UK for affected employees. That is a lesson to all of us to not dismiss the internal audit team.

In our scenario, the insurer suffers a loss of over £200 million on a 1 billion motor portfolio. A large proportion of this relates to third-party compensations and fines. These are also figures with the greatest uncertainty. For example, fines on the GDPR could be as much as 4% of global annual turnover. For this insurer, which has a 10 billion global presence, the fine could be as high as £400 million. Our thought process behind the fine was based around historical reported fines which were up to £500,000 under the previous Data Protection Act, and an analysis from NCC security consultants, which reported that this could have been 100 times larger under GDPR.

In the last few months there have been unprecedented levels of fines issued to companies. The rhetoric from regulators has now turned to action and it is clear that they mean business.

Since July, the ICO has issued their maximum allowable fine of £500,000 to Facebook attacks under the previous Data Protection Act. It is possible that we might see the first post-GDPR fine by the end of the year. Fines have not just been limited to the ICO. Earlier this month, Tesco Bank suffered a £16 million fine from the FCA for a security vulnerability in 2016. The FCA was especially critical of failures by Tesco Bank's internal team to follow their own procedures corresponding to such an attack, which led to a 21 hour delay in notifying the specialist fraud team.

This is a lesson for us in the insurance sector as well. It is important to recognise that it is not enough just to have a process in place. We need to be able to demonstrate to regulators that this process is being followed.

Rory Egan mentioned that in the US Uber suffered $148 million fine, and that had some traits that we would hope would not appear in the insurance sector, which is hiding the data breach for a year and paying a ransom to ensure that the data was not misused. This shows the importance of culture and conduct of a firm and the regulator will focus on this heavily when issuing a fine.

The changing landscape of cyber is not only in relation to new threats but, as seen by these examples, it is in relation to the regulatory environment and the litigiousness of cybercrime victims. By using the CRO Forum framework, we can identify changes in specific cost components as we receive new information. The lesson here is to stick to the framework, document the rationale behind the selection and document the uncertainty for each cost component to help communicate this to users and then stay up-to-date with cyber events in the news to help with following iterations.

Miss Jaeger: Scenario 2 is about a ransomware attack on a life insurance company. They are a subsidiary of a FTSE 100 company with £3 billion worth of gross written premium and profits of £300 million. They not long ago started an IT transformation programme to modernise their IT systems. They outsourced this with a data service company. It is terrible timing for something to go wrong.

For this scenario, we decided upon a ransomware attack against a number of insurance companies from a tailored spear phishing campaign. For this insurer, the hackers got through and they go undetected for several months. This means that when they deliver the ransomware worm it infects almost all of the insurance company's systems, both production and back-up environments. The impact is that the operating systems become unavailable. Critical systems and services are inaccessible, and the data is encrypted.

If we look at the root causes: this is lack of staff awareness as at least one employee clicked on the phishing email. Could they have prevented the phishing emails getting through firewalls or detected the hackers sooner and then behaviours? No one reported anything as suspicious.

How did they respond? In this scenario, the management call a meeting and decide to pay a ransom. After investigations, they have managed to reduce this to £7½ million to cover only the critical systems held to ransom. As we are looking at a one in 200 the event, this did not release the data and there is precedence for this.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

Basically, they are doomed as everything grinds to a halt; workers cannot do their work and are sent home at a cost. A huge data restoration project commences. The biggest cost here is in relation to significant lapses of in-force policies with loss of confidence in the firm and the reputational impact starts to kick in.

The conclusion for this scenario is a risk capital charge of 6% of the company's total revenue. What is worse is if the data is identified as being stolen, then we would also need to add on the cost associated with the data breach.

What are the mitigation actions here? As well as segmenting critical systems, the three key areas for improvement are around detection and looking to model trends in the standard behaviours so that they can identify and investigate any anomalies; improving vulnerability scanning and carrying out penetration testing, for example.

The response, looking to establish a decision tree, say, if they have a ransomware attack; creating and testing their instant response plans and perhaps appointing an expert on retainer to be available if an incident does occur or look at things like cyber insurance. It is also about the ability to recover, so regularly testing these plans, feeding back and improving on their protocols and establishing a clear communication plan. All of these could help to ensure quicker response, quicker detection and faster recovery, which could then reduce any reputational damage.

Mr Mohamed: In scenario 3, we looked at a hack of an insurer's telematics device. In insurance, telematics devices are becoming more and more common as insurers look at ways to differentiate themselves in front of the consumer and help them mitigate their risk. Estimates suggest that there are 17 billion connected devices worldwide today and Internet of Things (IoT) Analytics are projecting 34 billion connected devices by 2025.

In this scenario, the hackers were able to break into the telematics device, which allows them to remotely access images from the camera as well as location data such as customer journeys and frequently visited locations such as their home and work address. The hackers published the photos and location data on line. This is a scary scenario. What is even more concerning is the plausibility of it as we have seen with the reported hacks of baby monitors earlier this year. Insurers need to be prepared to measure and mitigate the risk that come with using connected devices.

We have formed a timeline for this scenario to illustrate how long cyber incidents can take to resolve. In this case it was 5 years. In this example, all 500,000 devices were replaced, which made up the bulk of the cost. We estimated the event cost the insurer 18% of their annual premiums.

In this scenario, I will expand on how we used the NIST framework. The NIST framework helped us to identify risks and ways that we could mitigate them: the more disastrous the impact of the attack the more relevant the NIST section is to help mitigate the scenario.

For example, the identify and protect has significant impact on the chance of the event; and the respond, recover and protect section of the framework are the most important control areas when it comes to reducing cost.

Consider "protect": in the scenario attackers were able to gain access to the devices by guessing the user name and passwords for the device. In this case "adminadmin". It is common that IOT devices have simple default passwords, and often these do not get changed or there is no ability to have them changed. Recently there are signs of action being taken by governments: California announced that manufacturers of Internet-connected gadgets will need to equip their products with reasonable security features out of the box, which means unique passwords or a feature that allows users to pick their own passwords. However, this law will only go into effect in 2020 and this is only one location worldwide to address the problem.

Using the NIST framework we can look at access controls within the protect section which states that user and admin accounts should be well managed from creation to deletion. For this scenario, that could mean having strict controls over who has access to the device and this should

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

reduce the risk of unauthorised access, and to change any default user name and passwords in the device. Another protect control would be the use of protective technology, such as performing penetration testing to understand how the devices could be exploited and what can be achieved with those exploits.

Could cyber sit under a rule of thumb? The IQ Insurance Insider reported that in practice most insurers using an internal model had calculated their cyber risk to be between 8% and 12% of their total Sales Revenue. Can cyber operational risk have a benchmark and would operational risk benchmarks be sufficient to cover cyber risk alone? It is too early for cyber risk to have a general rule of thumb. However, a bottom-up approach to assessing cyber operational risk will not only help you come up with a rationalised answer, it will develop the understanding of cyber risk within the plan and help prepare the firm for when a cyber incident does occur. It is fundamental to involve different people across the firm.

Cyber is a multidisciplinary problem, and introducing different perspectives can help challenge views and scenarios to get a better outcome. Our own scenarios changed as a result of consultation with other actuaries who joined the working party halfway through as well as when we ran them past IT specialists and cyber security experts. The CRO Forum and NIST framework can really help with communication between the different disciplines. Varying the scenarios as much as possible will help broaden thinking, as it did for us.

The type and size of insurer matters and account needs to be taken of the specifics of your entity: type of insurer; number of records held; how joined up the data systems are; the sensitivity of that data; all the way through to how your company interacts with their customers, what level of training the staff receive and the success of phishing tests. In short, what does the company's digital footprint look like, and what is the cyber resilience of the firm? All of these things will affect the chances of a successful attack and the cost of them.

Finally, we chose extreme but plausible scenarios and looked at setting the cost of a one in 200 return period. However, we recommend starting the scenarios and costs at lower return periods and then extrapolating to the one in 200.

The main learning outcomes that arose from this exercise were how to build a scenario structure and taxonomy, and a cost structure. This is important, especially when communicating with IT experts and with other actuaries. Think about the threat actors and vectors. You may not think it important to understand how you are being attacked, but that can have a significant impact on the cost and the sophistication of the attack. It is important to consult cyber and IT experts. When it comes to mitigation, it is equally important to involve as many departments as you can. For example, issuing cyber training would be via the HR department, and it is fundamental that they understand the benefit of this.

The environment changes rapidly. Writing this paper over the last year and a half, the environment changed multiple times. Scenarios that were considered extreme became obsolete, even in the short timeframe that we had.

Miss Jaeger: Firms need to take cyber security seriously. What we have demonstrated in these three scenarios is that although it is still a lot of uncertainty, the costs could be significant. Further, if there is a data breach, it is not yet clear what the potential fines could be under GDPR and there is that hanging question of whether these will be set high to make an example.

When it comes to the framework, our view is that NIST framework is a really good starting point for assessing key operational risks, but there are other frameworks which could be suitable for your company.

Regardless of what framework you use, the figures are likely to be highly subjective so it is important to create something that captures the specific circumstances of your company and draw upon the experience and knowledge from a range of experts that you have available.

Finally, it is important to have consistency and transparency. That is for communication purposes, but also to keep pace with the future developments. It is highly likely that the cost

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

calculations will change over time as threats evolve and as knowledge deepens, and it is useful for you to have an approach that can keep adapting to this.

Mr Egan: In terms of what happens next, we will be interested in your thoughts, but we have a few ideas to draw out based on the challenges that we had when putting the paper together. It is largely around sharing of knowledge, experience and data. We hope that this paper is the first step in the right direction. We would be interested to hear how you get on drawing methods and processes from it. I would be dismayed if you read it and speculated "Ah! The number is 6%. That seems fine. Job done!" That is not what it is about at all. It is about the framework, and applying that to your specific circumstances.

Data has been mentioned: that was a challenge for us, to find the right data points out there in the public domain to quantify the cost components. There have been some efforts in this regard. We followed on what the CRO Forum had already done and they in fact did a pilot study with participating insurers and operational risk people at those insurers to share data on cyber incidents and near-misses, those organisations that had suffered, and sharing them in an anonymised way.

The pilot was not extended beyond the pilot phase. There were legitimate concerns from the participants around the data that was being used: what is in it for me if I share it? Is it worth the cost/benefit analysis of that? From our perspective, as researchers, we just want the data. More data would enable us to come up with more realistic and plausible cost estimates and scenarios.

My understanding is that the Geneva Association has picked up the baton on this topic and is currently investigating whether there is scope for sharing of incidents, although they are focusing more on insured losses rather than operational risk events that insurers suffer. Equally, that will be helpful data. We would like to see more data, more knowledge sharing, more experience sharing.

What would you like us to do next? We could do more scenarios. We could look at pensions, for example. Then, more broadly, what should we as actuaries be doing? Do we need more specialised education, guidance? What should regulators be doing? Do we need stricter rules around cyber security? We have a range of opinions on that. One thing is for sure: it needs more attention, more people looking at the problem, more people getting schooled, learning about cyber risk and getting better at applying these techniques to determining how much of a risk it is. It is a risk where mitigation exists, and we are trying to find better ways to quantify the effectiveness of various risk mitigation activities.

Mr J. V. Moncaster, F.I.A. (First speaker from the floor): When quantifying the losses, how did you take into account the controls that were in place and did that then mean that you assessed the losses as less if there were more robust controls?

Mr Egan: We were working at the one in 200 where you assume that even with the controls you have in place they are not working or they are not implemented properly. We had some discussions about how to start to quantify the effectiveness of various controls and decided not to quantify them. We do list the most important controls as we see it for each scenario but do not give a quantification of how much it reduces it.

That is appropriate at least to assume that they are not going to work at that extreme return period, the one in 200. That was a huge challenge for us, and it was a challenge enough already to come up with the worst-case scenario.

Mr K. M. A. Tawfik, F.I.A.: Have you considered the possible recoveries from insurance contracts on the scenarios?

Mr Egan: That falls into risk mitigation to reduce the cost, which we have not looked at specifically.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

On the topic of cyber insurance, my view is that it is not a replacement for good risk mitigation. From an insurance perspective, why would they want to insure anybody that has terrible security controls? It is one of the tools. If you suffer an event, you can get some support on that.

We did not put quantification on it. But if you are able to buy cyber insurance, I would highly recommend that.

Mr M. L. Pearlman, F.I.A.: It sounds as if you already have two avenues for further research. I was wondering whether you were surprised at the numbers that came out of your research. Were they surprisingly high or low?

Mr Mohamed: On the data breach scenario, we were relatively surprised. However, there have been a number of data breaches in the past, not just in insurance but across all industries. We did have a few more data points. Those, as a proportion of the revenue, gave us an indication of where we thought the cost estimates would likely land.

Again, we want to emphasise the range of uncertainty in our estimates. When GDPR truly kicks in, they may be issuing fines of sizes that we have not yet seen before. It is definitely one that we should continually keep abreast of.

Mr Egan: If I may put it back to you, are you surprised that they are so high or low?

Mr Pearlman: The actual GDPR fines, the fine element of those numbers, were not that large compared to the other part of the costs involved. In a sense, the level of fining is irrelevant to the other factors in the calculation.

Mr Egan: We could have gone a lot further. Ramiz outlined the rationale on the fines. But we could have slapped the full 4% of global revenue and come up with a much scarier number and possibly would have been justified in doing so at a one in 200, which indicates the range of plausible and acceptable approaches and outcomes that one can come up with.

Mr Mohamed: One aspect is the size of a company: we could see much higher losses with the proportion of revenue on smaller companies than on the larger companies. Some of these costs are fairly fixed. You may need to hire a PR firm and lawyers to deal with incidents, and that is much higher proportion of revenue if you are a smaller firm.

Mr D. J. Hindley, F.I.A.: Rather than trying to share data on actual cyber events have you considered talking to insurers to see whether they would share with you some of the work that I know many of them has already done in this area, trying to quantify the impact on them of cyber risk events?

It will be at varying levels of maturity, but have you perhaps considered whether you could do a survey to try to refine or validate each scenario that the working party has done?

Mr Egan: Possibly, that could be a good avenue for us to explore now that we have the paper out there.

The Chair: We were looking at some of the scenarios and talking through how you came to those conclusions: you mentioned the importance of seeking external advice implying that there is no obvious answer. How do you achieve that balance, or weighting, between those in the business who may know the weaknesses and strengths, and an external consultant who perhaps will come in and have different views?

Mr Mohamed: An important question and one we grappled with. Having Ryan Rubin as part of the working party, who is a cyber expert, was invaluable, as he was able to determine how

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

realistic the scenarios were. He was able to give us a storyline that made sense, and we were able to understand. He deals directly with companies on mitigating their risks, on developing table top exercises.

This exercise cannot be done without cyber experts, whether external or internal. It is important to have both views if you can just have the comparison, especially given the levels of uncertainty in this field so having things that you can compare is certainly beneficial.

The Chair: A related question is if you follow this framework and work through the scenarios, then how do you maintain that? Everything we know today suggests that the world keeps changing, so how do you keep on top of something like this?

Mr Egan: Regular reviews: when we carried out the life insurance ransomware scenario, we initially proposed a targeted attack because that seemed to be the most extreme, but then NotPetya happened and we saw companies who are not specifically targeted suffering hundreds of millions in loss. So that changed our thinking.

The way to review appears to be by a step-by-step process. You can start with a top-down view and review it at a top-down level. Do we still have the right scenarios? Like a ranking system, almost, which helps you to decide whether or not you need to refresh and throw one out and bring a new one in or not.

Mr Egan: We go into detail when we quantify about the specific narratives and the specific cost items. Inevitably, the actual events that will happen will be different. It is not about getting the exact narrative right, but something that is representative of the type of risks that are faced.

Mr J. S. Halberstadt, F.I.A.: Given that the Allianz survey showed that this was the top thing keeping risk managers awake at night, how relevant is a one in 200 scenario treatment? Should it be taken out of that framework?

Mr Egan: That is a good question. Focusing on the most extreme, you are going to miss the smaller ones. To give you an example, at Munich Re we measure risk at the one in 1,000-year return period. So we are going well beyond the one in 200 in the other direction, as we need to look at the whole spectrum. You have to start somewhere and we started at the one in 200.

Mr R. M. Hill, F.I.A.: In scenario 2, the decision to pay the ransom, can you provide some information about that? Is that the worst-case scenario? Are people who have paid ransoms been treated differently? Is there an ethical issue there?

Mr Egan: There would be an ethical issue that would need to be considered by the company in question. Is it something they are comfortable doing, enabling this ransom business to continue? The alternative could be pretty bad if you do not pay the ransom and what that means to your business as a going concern. It is a realistic option that unfortunately needs to be considered.

In terms of the magnitude of the ransom amount that we assumed, it was higher than what we have seen so far. The largest one that we have seen is probably a bit lower (around 5 or 6 million). We are assuming that a £7.5 million ransom is paid.

But the problem in our scenario is that by paying the ransom you do not get your data back. It has just gone, deleted. It is a kind of double whammy – the one in 200 – and there is precedent for that from the NotPetya attack, for example, which appeared to be a ransom attack on the surface. If you tried to pay, you realised there was nothing behind it and they had just wiped all your data. So, we are going very much at the extreme end, but there is precedent for that.

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

On that topic, there is an organisation called nomoreransom.org that you can go to which you can sign up to as an organisation to try collectively to stop these ransomware attacks by refusing to pay the ransoms.

We have to be realistic about what is happening. These ransoms are being paid because the alternative is too scary to consider for many organisations.

Mr Moncaster: What threshold are we looking at whether it is the one in 200 or not? You talked about extrapolating from a more likely scenario. Could you talk through the approach that you envisage that people would use towards extrapolating?

Mr Mohamed: The traditional method on the underwriting risk side is ask the underwriter for the one in 5, the one in 10, the one in 25 and the one in 50 and then use that as data points to fit one of four distributions to it. This could be a method, especially if a lower return period is more tangible in terms of loss estimates and the data behind it.

As we get more losses, or as the risk matures, it is very possible that it will end up being a risk similar to how we treat other risks. That is certainly one way of tackling it.

Mr Egan: I would also look at it from the other way, that is consider the worst-case scenario and then imagine most of your controls that are expected to work do not work. Maybe that is not a one in 200; maybe it is a one in 100, one in 500. It is going to be an extreme number and that is more of what we were considering.

Mr A. Kaye, F.I.A.: Most firms outsource. How do you factor that into your modelling?

Mr Kaye: The fact that most firms outsource to a greater or lesser extent, so the attack could be impacting the outsource supplier, not you directly, means you do not have control over remedial actions and are not in direct control.

Mr Mohamed: This is covered in our third scenario, a telematics device, which was manufactured externally. The outsource risk is addressed in the latest NIST update with a specifically beefed-up section around supply chain risk. The concern these days is everything is connected to everything else. So even if you know your suppliers, you do not necessarily know their practices and who is connected to them. That chain can go on and on, so absolutely a scenario around supply chain risk is certainly one to develop.

Mr Kaye: A slightly different question: your definition of cyber risk includes equipment failure. So one would imagine the approach to that for a hardware failure, for example, a backup system not working, we see a number of examples in the marketplace, would be rather different to the ones that you have been addressing until now.

Mr Mohamed: Yes, for example, the Spectre/Meltdown last year was a potentially aggregating hardware failure event that would certainly have severe impact on some companies. I agree with you that hardware failure is as important as people failure or software failure.

Mr R. Laux: One of the challenges we face is that everything is connected to everything else. Yet our IT professionals where I have worked are always lamenting the fact that systems do not talk to each other. That is viewed as a bad thing. But it sounds from your perspective systems that do not talk to each other would be a good thing by limiting the damage of an attack. What do you think?

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

Mr Mohamed: As insurers we may benefit from being behind the curve because if one system falls over, the rest is fine. I cannot even get all the data I need in my own company, let alone someone externally getting it. No doubt there is a benefit to being disaggregated in that respect.

Miss A. Yeap, F.I.A.: Do you look at the credit monitoring services, the cost of offering the service to the customers where their status has been breached, and do you look at the proportions of people who will take up that insurance?

Mr Mohamed: Yes, we have included that in the data breach scenario. It is one of the typical first-party losses that companies are to address. We have found that those costs were a very small proportion of the total loss, which was mainly driven by compensation or reputational damage or maybe third-party-type liability losses, more so than the cost of credit monitoring or notification costs.

Mr Egan: Yes, but that is probably one where you need to keep in mind we are working at the one in 200-year return period. The idea behind providing credit monitoring or things like that is you are trying to prevent them from suing you. So you are giving them stuff proactively so that they do not say: "You have done this bad stuff to my data and I am going to sue you."

It is something that you would see as part of a typical response for a data breach. First-party costs spent now to avoid possible scarier losses on the third-party side later on. But in the one in 200 we are looking for the worst-case scenario, so the third-party costs are bigger.

The Chair: I was wondering, a little bit broader, as we have been discussing today about the scenarios and things that are really important to try to capture cyber risk, do we think that this is where we are going to be for quite a long time?

It seems from conversations about data we are all suggesting that we are a long way from having anything where we have a magical model. Could we ever get to that state? Things are constantly changing. Is all that data not going to be of any use because the world will not look the same next year as it did last year? Is it likely that this is going to form a long-term view of how we look at cyber?

Mr Mohamed: The benefit of using scenario modelling is it is tangible, it is easy to communicate, so laying out a storyline is always a good way of getting people to understand what has happened and what the impacts are.

I agree that there may be other ways of doing it in the future, even in an evolving landscape. For example, some colleagues of mine went to meet Microsoft. They have some incredibly advanced techniques to understand where there may be points of failure in their systems.

If we had access to that sort of level of information, as more companies migrate to the Cloud, as more companies outsource their IT, it is very possible that we could come up with much more sophisticated views in the long run. Scenario-based modelling will always have a place because of the communication element.

Mr Egan (closing the discussion): If anybody has any suggestions, based on questions here, about what should happen next, we would be interested in your opinions. Do we need more regulation? I do not think that we have a great deal in the UK right now other than a keen interest from the regulators.

What should we as actuaries be doing? How do we improve overall cyber risk awareness, getting more people aware, reducing those weak links that I talked about within each company? Addressing the challenges we face around finding data to quantify cyber exposure. If there are any questions or comments on that, I would happy to hear them.

Equally, now we have finished this work stream around cyber operational risk scenarios for insurers, we are now looking for our next project. So if you would like us to look into any topics

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press

in particular under the broad umbrella of cyber risk, where the actuarial profession can show value, then we would love to hear from you.

Mr J. D. Tapper: Your work is based on operational risk for insurers. But what is the appetite within the group for researching silent cyber risk? That relates to claims to insurers that are not from direct cyber policies.

Mr Egan: With silent cyber risk (for people who are not familiar with that), we have the cyber insurance market on one hand, which are products specifically designed to pay out in the event of cyber events. You then have the inherent cyber risk that is sitting in property and casualty policies that are covering fire and liabilities and things which were not designed with cyber attacks in mind per se, but are going to be exposed.

There appears to be a lack of understanding and research on the topic of how big the problem is; going through line by line looking at property, fire insurance, is there a big cyber potential there compared to something on the liabilities side?

We have not looked at it yet as a working party. It is definitely one of the topics we would consider strongly for our next piece of work. It is a big issue.

There has been some work on it. Willis Towers Watson did a survey to try to gauge how big a problem we are talking about. They asked people in the insurance industry line by line what proportion of the premium for property, workers comp, other casualty lines, for marine, what proportion of that do we think would we need, or additional amount, would we need to cover the silent cyber risk. There have been a few things but we are only scratching the surface there. That is one potential area we would consider looking into.

Another one is accumulation risk, whether we as actuaries, and volunteer actuaries, are the right people to do all these things. I do not know. These are the topics that we are facing.

Accumulation risk on cyber is real. Munich Re, Swiss Re, others, are having these debates about whether you could insure the Internet as in if the Internet goes down in a country. Is that even possible? Would it be long-lasting? Would it be 5 minutes? Would it be an hour? Could it be days?

Are these sorts of risks even insurable, and if they are insurable can they be quantified, modelled, et cetera? So, silent cyber and cyber accumulation risk are two possible topics for us to investigate further.

The Chair: Thank you. It remains for me to say thank you to the audience for coming and contributing towards the discussions and questions that we have had and thank you to our speakers, not only for putting in the time today, but also for the research. Members of the working party have been working for years in some cases on their research and the culmination that we have today.

Just to conclude, if you do have any more comments or questions that you want to direct to us afterwards, please do so. A copy of the paper and slides are available on the Institute website if you wanted to pull off that information.

Cite this article: Cyber operational risk scenarios for insurance companies. *British Actuarial Journal*, 24. https://doi.org/10.1017/S1357321719000138

https://doi.org/10.1017/S1357321719000138 Published online by Cambridge University Press