Lessons on Catastrophe:
Differences and Similarities between Cyber and Other Forms of Risk

Authors
Rob Champion, MSc (Royal Holloway, 2019)
Carlos Cid, ISG, Royal Holloway

Abstract

The field of Cyber Insurance is still in its infancy but has already shown significant growth with plenty of evidence for further expansion. However, a lack of past information and some idiosyncrasies make pricing difficult as well as potentially amplifying risk exposure.

This article summarises high level findings from a practical model that could be used in lieu of actuarial data. The model may be refined in the future as historic datasets become available. This practical model shows that Cyber Insurance risks pose significantly elevated likelihood and impact when compared with other forms of risk which are more independent. Higher premiums will be a natural consequence to insulate from the associated downside.

There are therefore strong incentives for Insureds to improve event independence, for example through hardening. Insurers, on the other hand, can protect themselves from extreme events by rejecting certain risks with cover limits, as they do already, or they may choose to transfer the more extreme risks via commercial structures and Insurance Linked Securities, particularly Catastrophe Bonds. [a]

[a] This article is published online by Computer Weekly as part of the 2020 Royal Holloway information security thesis series `https://www.computerweekly.com/ehandbook/Royal-Holloway-Lessons-on-catastrophe-differences-and-similarities-between-cyber-and-other-risks`. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full thesis is published on the ISG's website at `https://www.royalholloway.ac.uk/research-and-teaching/departments-and-schools/information-security/research/explore-our-research/isg-technical-reports/`.

Introduction

Cyber security is one of the most pre-eminent risks of current times. Whether it's the theft of consumer data by the likes of MageCart group, corporate disruption by ransomware such as WannaCry, or highly sophisticated attacks by APTs (advanced persistent threats), today's headlines are drawing attention to risks in the digital world.

At the same time, demand for risk management products is rising too. Assuming the risk cannot be avoided, and that acceptance is too perilous, there are two remaining alternatives. The first, mitigation, is what's powering the growth in the cyber security sector in recent years. The second, transference, is the driving force behind Cyber Insurance, as it allows realised costs from risk to be shared with a third party, the Insurer.

Given the recent growth trend of this sector, with a Compound Annual Growth Rate (CAGR) in the double digits for both actual and forecast, there are several important questions that can be asked:

1. What differences set Cyber Insurance apart from other forms of insurance?
2. How might Insurers model the idiosyncrasies of Cyber Insurance?
3. What measures could Cyber Insurance buyers (the "Insured") implement to reduce the cost?

The Insured's Perspective of Cyber Insurance

When comparing Cyber Insurance with other forms of insurance, it's helpful to have a framework. From the Insured's perspective, one such option is the "Six Cs"[1], as follows.

Coverage: the quantitative limits and qualitative exclusions of the policy.

Cost: comprised of fixed periodic payments known as the Premium, and the threshold of costs before the insurance pays out, referred to as the Deductible (US) or Excess (UK).

Capacity: the Insurer's ability to handle any loss events that occur to the Insured.

Capability: offerings beyond the insurance itself, such as consulting or forensics.

Claims: how loss events will be handled.

Compliance: adherence to legal requirements of both the Insured and the policy itself.

According to a Mactavish report[2], the most concerning issues for those who took out insurance were Coverage, closely followed by Cost. The only other category mentioned in the report is concern over the Insurer's Capacity to cover the full value of arising claims.

Fittingly, these categories are also where digital risk is an outlier. Whether it's determining if NotPetya falls under Force Majeure as an act of cyberwar, or the finding that premium-to-cover ratios for Cyber Insurance are 3 times higher than found in general liability insurance, Cyber Insurance doesn't quite fit the mould of more established forms of insurance. Thus far, Capacity issues have not arisen but as we shall see, Cyber Insurance does pose challenges to Insurers on this front too.

The Insurer's Perspective

A similar framework by Berliner sets out categories for an Insurer, from which it is apparent that the Insureds' perspective is reflected in the first two categories which comprise Market Risk, though the source of the anomalies for Cyber Insurance lies in the Actuarial categories, on which pricing is based.

	Insurability Criteria	Risk Type
1	Insurance Cover Limits	Market
2	Insurance Premium magnitude	Market
3	Information Asymmetry (Moral Hazard & Adverse Selection)	Actuarial
4	Randomness of Loss Occurrence	Actuarial
5	Average Loss per Occurrence	Actuarial
6	Average Frequency of Occurrence	Actuarial
7	Maximum Possible Loss	Actuarial
8	Public Policy	Societal
9	Legal Restrictions	Societal

Significant academic effort has been directed at these actuarial concerns. The lack of information was the first issue to be highlighted and remains a significant problem for analysis of the risks to Insurers in a sector where they must rely on the information provided by the customer. The complexity of cyber risk makes it much more difficult for the Insurer to get the complete picture, as well as contributing to two behaviours that further increase risk:

[1] Hopkin, P. (2017). Fundamentals of Risk Management. London: Kogan Page.

[2] https://www.mactavishgroup.com/wp-content/uploads/2018/11/Mactavish-Cyber-Risk-Insurance-Report-November-2018.pdf

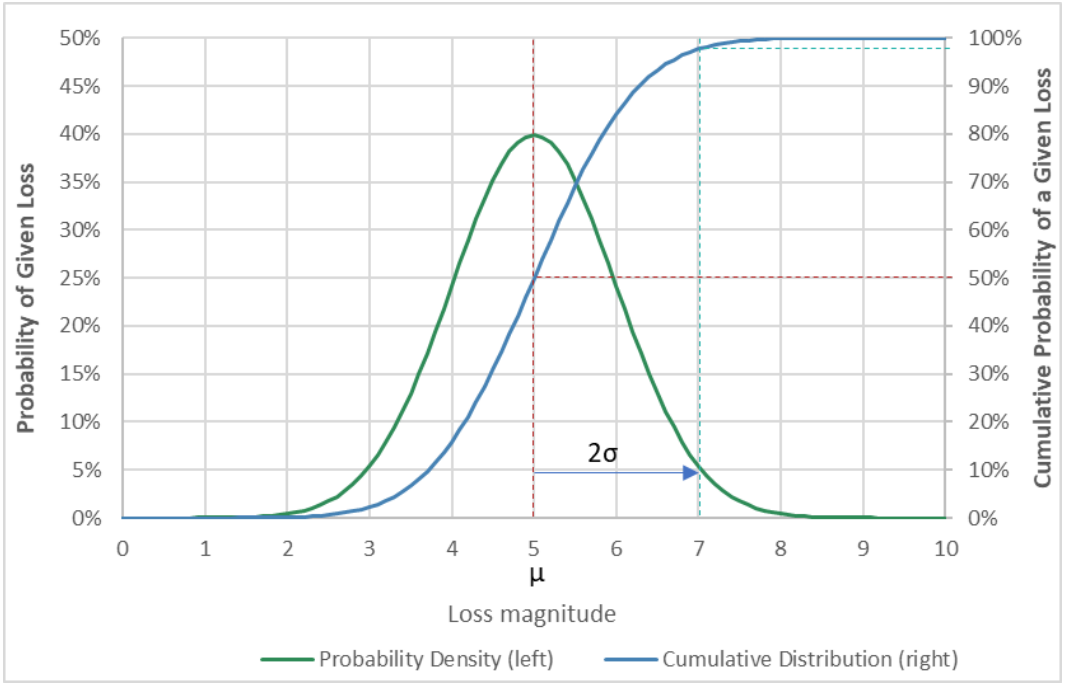

Figure 1: A fictional risk that is normally distributed with a mean cost of 5 units and a standard deviation of 1.

- Moral hazard, when the Insured behaves more riskily than otherwise, under the impression that they are protected.
- Adverse selection, when high-risk parties purchase insurance, while low-risk ones don't.

This has since been joined by concerns over *randomness of occurrence* with implications on Insurers capacity. In most forms of insurance, incidents are independent. A car accident in one country is unlikely to be directly connected to another in a separate country. Cyber incidents, however, are much more interdependent and an incident affecting one Insured can certainly cause problems for another, such as via the supply chain.

This all contributes to the high premiums and tough cover limits mentioned previously. But how might this be all be quantified?

Modelling Cyber Insurance

Before going into detail on a potential model it is worth briefly outlining an approach used to price insurance that is both simple and avoids intrinsic assumptions of risk tolerance, called the Standard Deviation Premium Principle.

Figure 1 shows a probability graph illustrating a fictional risk that is normally distributed with a mean cost of 5 units and a standard deviation of 1. The absolute minimum that a premium could be set in this example would be at the mean, or 5 units, but that also means that there's a 50% chance the loss will exceed the premium and the Insurer will be unable to pay (known as the Risk of Ruin).

Perhaps a more sensible approach for the longevity of the Insurer is to charge the expected value plus two standard deviations ("$\mu + 2\sigma$", or 7 units), which means that there's now only a $\approx$ 2.5% risk of ruin. This would greatly improve the Insurer's Capacity, but at the expense of the policy Cost. Thus, the magnitude of the premium would be a function of three figures:

- average expected loss, μ;
- standard deviation, σ, which is a measure of the volatility or risk of a policy;

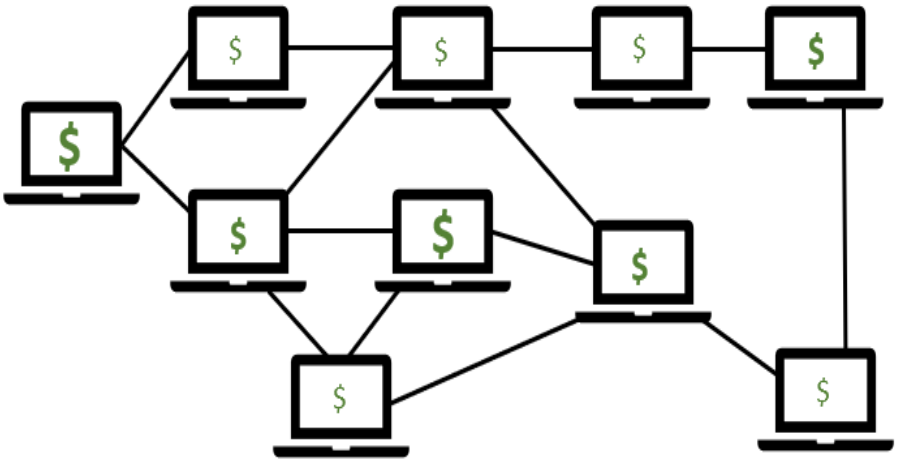

Figure 2: Network of systems.

- multiple of standard deviations, in this case 2 but determined by the Insurer.

Thus, we can now determine the cost of a policy based on the probability and costs generated by a model.

The first step is the model structure, in this case a Monte Carlo simulation running iterations of a year in the life of a network. This network is comprised of “nodes” to represent systems and “edges” that link one system with its networked neighbour, illustrated in Figure 2. The systems may be individual computers, or an entire contained network.

Finding a suitable common variable with which to represent impacts is difficult, but in this case, revenue was chosen. While there are pros and cons of all variables that could be used, revenue is widely present, with notable exceptions such as start-ups, generally scales with an organisation and is easily relatable. Impacting factors will therefore result in disruption to this value before recovery can occur.

The next step is to determine how impacts can arise and typically fall under two mechanisms:

Random selection: independent impacts exemplified by component failure or untargeted phishing compromise.

Propagation: interdependent impacts arising from the spread of a malady from one affected node to another, exemplified by malware and traversal of a compromised network.

The final step is to model a propagation mechanism. This property makes cyber risk unusual although not unique – this approach was pioneered previously in the study of disease transmission and epidemiology.

Unlike prior models, which simply use probabilities to determine infection, this one uses a more adversarial approach. An Attacker can use one or more exploits in its efforts to compromise a Defender, who in turn may have any number of vulnerabilities. If there is a match, and the Attacker’s exploit has a higher version than the Defender’s vulnerability (i.e. the Defender is unpatched) then the contagion spreads (see Figure 3).

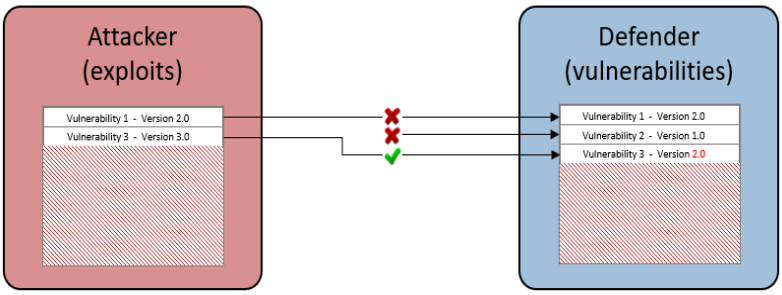

Figure 3: An adversarial approach.

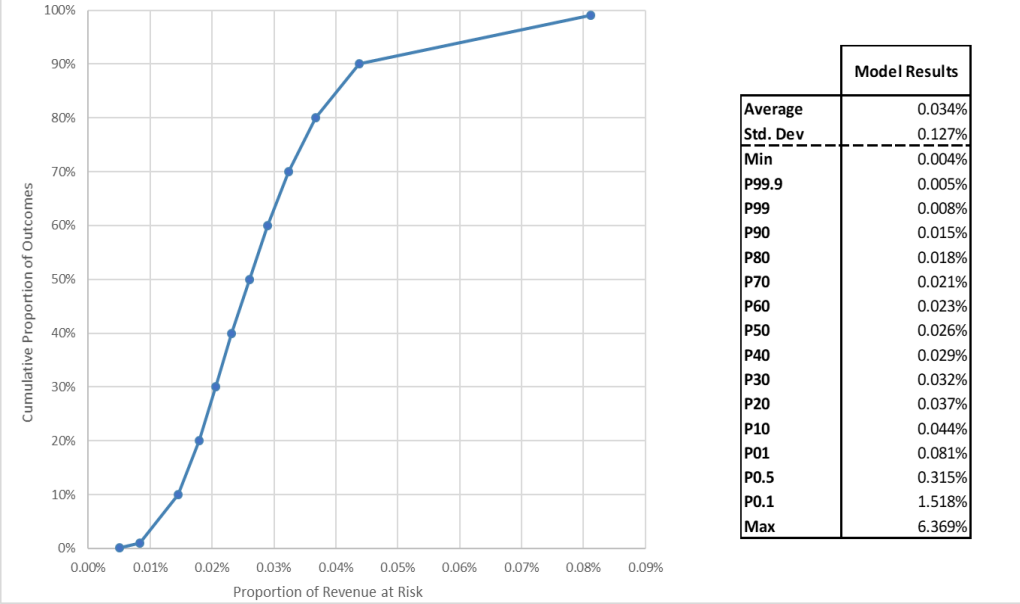

	Model Results
Average	0.034%
Std. Dev	0.127%
Min	0.004%
P99.9	0.005%
P99	0.008%
P90	0.015%
P80	0.018%
P70	0.021%
P60	0.023%
P50	0.026%
P40	0.029%
P30	0.032%
P20	0.037%
P10	0.044%
P01	0.081%
P0.5	0.315%
P0.1	1.518%
Max	6.369%

Figure 4: Results of model.

Results

The results of this model (Figure 4) show reasonable proportions of revenue at risk with a less than 1-in-100 chance of an impact exceeding 0.1% of a company's revenues. At such a level, losses would be insurable, as well as bearable (if undesirable) for an uninsured business.

What is omitted in this graph is the top 1 percentile. While to many this may appear to be "outliers", this is far from the case for Insurers. In fact, the Solvency Capital Requirement (SCR) for insurers is set at the 99.5% confidence interval (i.e. the P0.5), making these extreme outcomes of great interest. It is in the final percentile that the severity of risk in Cyber Insurance, both in terms of magnitude and likelihood, becomes apparent.

The sequence of charts in Figure 5 shows the severity of the skew beyond the 99% of outcomes and is important for two reasons: First, the long tail is less likely to be coverable by a business' own means, either through self-insurance or free cash flow. Consequently, insurance becomes more important for business survival in the event of such a major outcome. Second, Insurers need to ensure they are adequately protected from such outcomes too. This would be achieved through greater safety capital and/or higher premiums which, of course, makes the policies more expensive to the Insured.

When certain variables are isolated, the effect of mitigations applied in the model can be measured.

The simple, common Information Security recommendations of reducing vulnerabilities (i.e. patching) and network hardening both clearly dominate the base case. Network Hardening showed a 15%

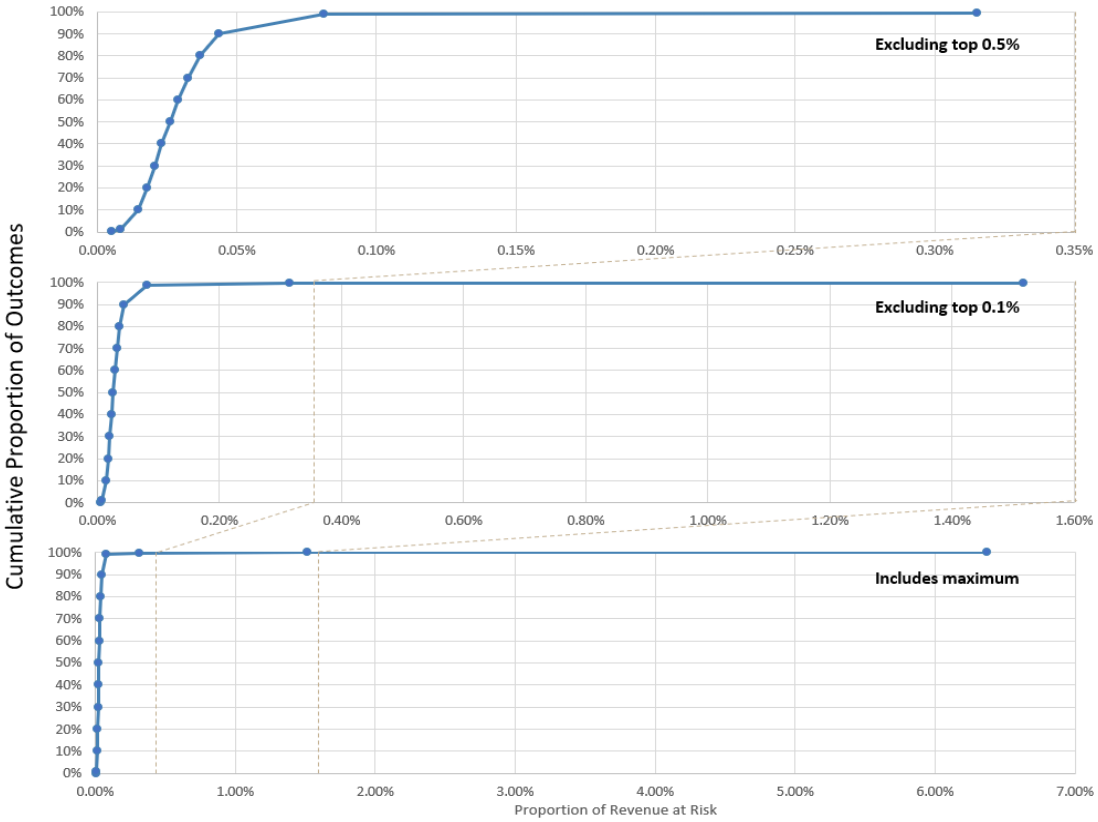

Figure 5: Severity of the skew beyond the 99% of outcome.

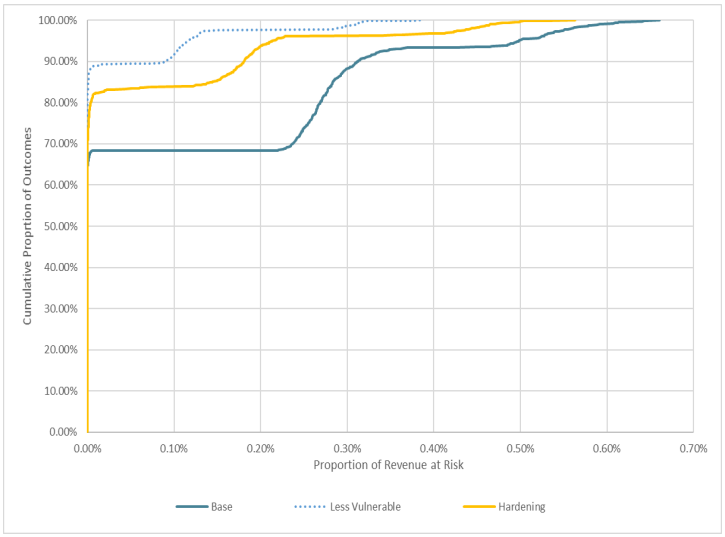

Figure 6: Outcomes.

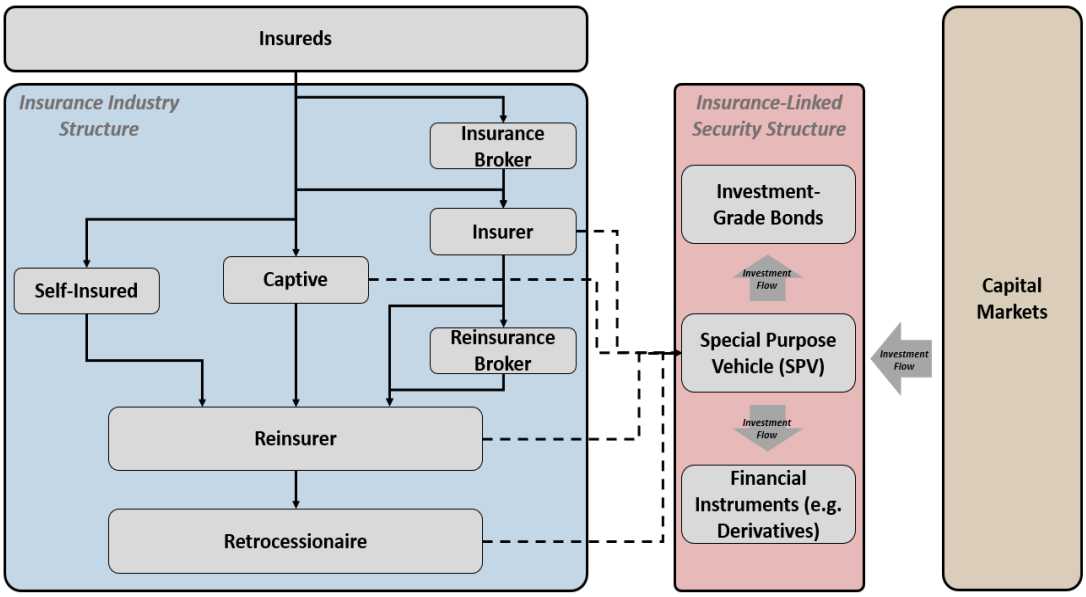

Figure 7: Catastrophe Bonds.

improvement in expected losses. Meanwhile patching improvements, roughly halving vulnerabilities at any given point in time, showed a 43% reduction to the expected loss in cases where these factors play a role (Figure 6).

Interestingly, highly structured networks can be beneficial in this context, with certain caveats. Highly structured networks can limit propagation through the network provided an incident is identified in the process of spreading. However, this benefit may be difficult to capture as well as coming at the expense of increasing the fragility of the network to DoS attacks and failures. Depending on the business, DoS attacks are typically lower cost although they may also be highly visible.

Conclusion

Using what little information is available in this field, it is apparent that the same technology that's connecting the world can also magnify risk.

For Insureds, it shows the potential benefits to be had from certain Information Security investments by reducing ongoing premiums. Given the transmissible nature of some cyber threats, such investments can also serve a Common Good as reducing transmission vectors may help the broader society.

Both Insureds and Insurers can also choose to cap the risk transfer, leaving the risk of the most significant claims with the Insured. This would mitigate the extreme tail-end risk shown by cyber policies and reduce premiums while also preserving much of the functionality of the insurance in most cases. Insurers would also be able to limit their aggregated risk, too.

Insurers may also diversify their portfolio through other non-cyber policies. This would mitigate exposure by sourcing much more independent sources of risk.

One final mitigation may be achieved with the use of Catastrophe Bonds, a type of Insurance Linked Security that bridges the gap between the Insurance Industry and the wider Capital Markets, per the diagram in Figure 7.

Insurers of all levels can transfer the most severe risk into a product available for investment from the wider Capital Markets. Investors would be entitled to receive periodic income sourced from the premiums in exchange for providing a lump sum that may be eroded in extreme events. This is akin to purchasing Corporate Debt and, via Credit Ratings, may offer another measurement with which to

gauge risk.

So, while Cyber Insurance poses challenges when compared to more conventional forms of insurance, we can use learnings from sectors as diverse as healthcare and corporate finance to help inform the models used to analyse the market.

As George Box said, “All models are wrong, but some are useful”. Hopefully this one manages to achieve the latter.

Biographies

Rob Champion is a CFA Charterholder with over a decade of experience in financial analysis and operational risk management in the Oil & Gas sector. He now runs FilaContexta, a financial modelling consultancy, using his studies for an Information Security MSc to help develop models of more complicated sectors.

Carlos Cid is a Professor in the Information Security Group at Royal Holloway University of London, and a Visiting Research Professor at Simula UiB, Bergen, Norway. His main research interests are cryptography and cyber economics. He has a PhD in Pure Mathematics from the University of Brasilia, Brazil (1999), and before joining Royal Holloway in 2003, held academic and industry roles in Brazil, Germany and Ireland. Carlos was the founding director of Royal Holloway’s Centre for Doctoral Training in Cyber Security, a post he held between 2013 and 2017.

Series editor: S.- L. Ng