Security Economics Knowledge Guide
Issue 1.0.0

Tyler Moore | University of Tulsa

EDITOR
Yulia Cherdantseva | Cardiff University

REVIEWERS
Ross Anderson | University of Cambridge
Daniel Arce | University of Texas at Dallas
Rainer Böhme | University of Innsbruck
Jason Nurse | University of Kent

COPYRIGHT

© Crown Copyright, The National Cyber Security Centre 2024. This information is licensed under the Open Government Licence v3.0. To view this licence, visit:
http://www.nationalarchives.gov.uk/doc/open-government-licence/ OGL

When you use this information under the Open Government Licence, you should include the following attribution: CyBOK © Crown Copyright, The National Cyber Security Centre 2024, licensed under the Open Government Licence: **http://www.nationalarchives.gov.uk/doc/open-government-licence/**.

The CyBOK project would like to understand how the CyBOK is being used and its uptake. The project would like organisations using, or intending to use, CyBOK for the purposes of education, training, course development, professional development etc. to contact it at **contact@cybok.org** to let the project know how they are using CyBOK.

1 INTRODUCTION

The subdiscipline of *security economics* was instigated by Ross Anderson in 2001 [4], with the first annual Workshop on the Economics of Information Security held in 2002. This knowledge guide does not seek to summarise the entire field in a few short pages, an impossible task to be sure. Instead, the goal is to introduce the reader to some of the most impactful ways in which economics has helped to shed light on cybersecurity problems and frame solutions that blend private and public action. The guide focuses on the organisational, rather than individual, perspective, which is where the majority of scholarly activity has focused.

Section 2 describes canonical security failures from an economic perspective. Section 3 describes key measurement challenges. Section 4 reviews firm-level approaches to improving cybersecurity while Section 5 discusses available public-policy options.

The CyBOK introduction identifies Cyber Security Economics as a cross-cutting theme [36]. We cross-reference relevant CyBOK knowledge areas throughout this guide for the interested reader to explore further.

2 SECURITY FAILURES

We know that security failures happen. Financial institutions leak account data on millions of customers. Municipal services shut down for weeks upon discovering that systems have been encrypted and rendered inaccessible until a ransom is paid. Nation-state hackers penetrate electricity grids, prompting fears of sudden shutdowns.

In each case, there are often perfectly understandable technical explanations for why the attacks succeeded, which vulnerabilities were exploited, and so on. By focusing on the details of how the attackers succeeded, it is easy to miss the more profound questions lurking beneath the surface: Why are so many systems vulnerable to attack? Why do some attacks such as ransomware continue to proliferate when common-sense defences could stop them from succeeding? Why is society less secure today even though we are spending far more on information security defences than in the past?

An economic perspective offers compelling answers to each of these questions, by explaining shortcomings that would otherwise seem baffling when only adopting computer science approaches.

2.1 Misaligned incentives

We begin by discussing misaligned incentives. Computer systems often fail because the organisation that ought to be in the best position to protect a resource does not always have the strongest incentive to do so. Anderson [4] discussed retail banking in the 1990s, which is an instructive historical example. While the environment of banking system security does not technically vary that much from country to country, the regulatory environment in which different banking systems operate does differ greatly. Comparing outcomes between environments can help to identify the role that incentives play. In the United States, banks have long been required to pay for ATM card fraud, credit card fraud and the like if a transaction is disputed [54]. Customers alert the disputed transaction to the bank and the bank reimburses the customer for the fraud, usually with no questions asked. However, not all countries work

that way. In the United Kingdom, for example, fraud regulations have historically favoured banks [3]. This has made it easier for the banks to disclaim liability for fraud and push the responsibility back onto the cardholder. For example, when a fraud dispute arises, the standard argument from UK banks would be that the customer was careless with their PIN and somehow must have disclosed it to the criminal.

So we have two countries with different regulatory environments – which country suffers more fraud per capita? The United Kingdom. The reason why goes back to the incentives at play. US banks have had to pay for disputed transactions, which meant that there was no getting out of it. This created a very strong incentive for US banks to invest in technologies that reduce fraud. Conversely, UK banks could blame customers for fraud in many circumstances. This dulled their incentive to invest in technologies to reduce fraud. Consequently, UK banks experienced much higher rates than observed in the United States [3]. This early example illuminated many researchers to the importance of understanding the circumstances in which a technology operates beyond the purely technical considerations.

Outside of banking, we see misaligned incentives appear in many other contexts. For instance, we see insecure devices and supply chains, as well as differing and competing interests in public versus private attitudes toward critical infrastructure protection. Let us take one deeper dive into the incentives involving critical infrastructure protection. Industrial control systems configure and operate many critical infrastructure systems, including power, pipelines and even wastewater treatment. The devices in control systems were not designed to be connected to the Internet because they have very limited security capabilities built in, and because they are often used to program other devices in the system. And yet a researcher found that thousands were globally accessible and available to all [34]. If the best practice is not to connect to the Internet, why do so many people do it anyway? We can find an answer by taking a closer look at the incentives at play among the various stakeholders.

Critical infrastructure operators definitely want to protect industrial control systems from attacks. However, there is a clear benefit to connecting systems to the Internet, in order to ease remote administration. Maintaining physical separation of IP-connected networks and industrial control systems inevitably drives up costs. This is compounded by the fact that historically there have been very few attacks on critical infrastructure, at least not many that have been publicly disclosed. Finally, if something actually does go wrong, the full cost of an attack is borne not just by the operator but also by society at large. Taken together, the rational decision for many critical infrastructure operators is to connect the devices online, despite the recommendations from best practice.

Consumers, meanwhile, value the reliability of critical infrastructure services and do not want them to be attacked. Yet the same time, consumers prefer low cost of service, and it is not feasible for consumers to distinguish between the security among critical infrastructure firms.

Governments, like consumers, value reliability of service, and they are especially wary of the potential of being blamed for an attack. Governments are definitely worried about cyber security, but they are not in a strong position; they do not have the budget to fund security in the private sector, nor do they often control it.

In the absence of any regulatory intervention, the stakeholder with the most control usually determines the outcome. In this case, it's clear that the incentives of critical infrastructure operators matter most because they're the ones who actually decide how to make their investments in security and make the trade-off between connectivity and security. Clearly, they are choosing to connect their systems to the Internet, which is what the observed experience

indicates.

2.2 Market failures

Economists use the term market failure to describe when the real world doesn't quite live up to the models of perfect competition. In fact, markets fail time and again in the same particular ways, so much so that several distinct categories have been identified by economists: monopoly, oligopoly, public goods, information asymmetries, and externalities.

In a monopoly, a good or service has only one provider. Under these circumstances, the monopolist affects prices by controlling supply. In oligopolies, only a few providers are available.

Most goods can be privately consumed. Material possessions such as cars and houses are assigned individual owners, and it is natural that no one else can consume the good at the same time. But certain goods cannot be privately consumed, such as investments in national defence or even the air we all breathe. Public goods behave differently than normal private goods in two key ways. First, public goods are nonexcludable: there is no practical way to prevent people who don't pay from consuming the good. Second, public goods are nonrival. When someone consumes the good, this does not limit others from also consuming it. National security is a classic public good. It is nonrival because an individual benefiting from the presence of a nation's military does not prevent others from experiencing the same benefits. It is nonexcludable because it is not feasible, let alone ethical, to exclude tax evaders from receiving protection even though they did not pay for it.

Sometimes, cybersecurity behaves as a public good. Cybersecurity is almost always nonrival – adopting security controls to protect a resource does not make it any harder for someone else to do the same. Whether a cybersecurity technology is excludable or not is often a decision left to the provider. For example, Microsoft used to only offer security updates to customers with verified software licenses, thereby excluding security benefits from those with unlicensed copies. They changed their policy and now offer security updates to all, because improved security often benefits others too (more on that when we discuss externalities below).

The last two market failures – information asymmetries and externalities – appear in many cybersecurity contexts.

Information asymmetries occur when one party in a transaction knows more than the other. A seller might know more than the buyer, or vice versa, about the quality of the good exchanged. This imbalance can lead to inefficient outcomes and security failures.

Used car markets offer the classic example of markets with asymmetric information [2]. Suppose a town has 20 similar cars for sale, so they have the same make model and mileage. In fact, 10 are high-quality "cherries" worth $2,000, while another 10 are low-quality "lemons" worth $1,000. However, it is impossible for buyers to distinguish the cherries from the lemons. Akerlof demonstrated that the market-clearing price in the presence of such asymmetric information is $1,000 [2]. This happens because buyers cannot distinguish high-quality cherries from low-quality lemons and so refuse to pay a premium for higher quality. Sellers, in turn, know this. So the only people who are actually willing to sell their cars for $1,000 are the owners of the lemons. Consequently, the market is flooded with low quality goods. This phenomenon is known as adverse selection.

Unfortunately, information asymmetries also plague cybersecurity markets [4]. First, the

market for secure software and products is a market for lemons. Vendors may believe that their software is secure. They may be marketed as such, but it is very hard for a buyer to credibly evaluate such claims. Consequently, buyers are unwilling to pay a premium for more secure software. Instead, they prioritise other features that they can measure the quality of, such as the user interface and price. Developers in turn put more effort into satisfying the qualities that can actually be observed. But this leads to a bad outcome because security is not emphasised as it should be.

We can see this in other areas besides software. Notably in security investment for firms more generally, where a firm needs to convince its customers that it respects their data. It can be very difficult to demonstrate good security. By contrast, poor security can be readily confirmed when a breach occurs.

The second broad area where we see information asymmetries arises from the lack of data relating to cybersecurity incidents. Companies prefer not to disclose when they suffer an attack. If they are not required to by law to do so, many firms simply will choose not to. While not disclosing may be a sensible strategy for an individual firm, this approach makes it very difficult to get a grip on the true nature and extent of cyber risks within the larger context of a lack of available evidence. Firms cannot easily estimate the probability that an incident will take place or what it might cost them. In other words, if nobody is talking about when an attack happens or what it costs them, then it is very difficult for other firms who have not yet been targeted to know what the real risk is. Thus, we have an information asymmetry that exists between firms. Even when firms do share that they have been attacked, they tend to avoid discussing the financial fallout. Thus, we end up with a lack of accurate information on losses, which makes it difficult to know if past investments have been effective. Measurement, as explained later, is a huge challenge.

What are the consequences of markets with asymmetric information? There are two classical outcomes: adverse selection and moral hazard. Adverse selection could happen in the cyber insurance market [46]. It is very difficult for an underwriter to discriminate between firms based on their operational security practices. Insecure firms are more likely to buy cyber insurance, which could trigger higher premiums and lower participation than otherwise. We also see adverse selection in the abuse of signalling devices, as discussed later.

The second outcome of asymmetric information is moral hazard. People may change their behaviour if they are given some kind of protection. For example, in auto insurance, people may drive more recklessly if they are fully insured and have a very low deductible because they know they will be covered. In practice, moral hazard has not been a significant issue in cybersecurity so far since insurance does not fully cover all harms that result from cyber attacks [28].

Moral hazard can also arise in situations other than insurance. For example, in a cloud environment, customers cannot observe the providers' security efforts while providers cannot always observe their customers' security practices. In both cases, each side may be tempted to shirk their responsibilities and operate insecurely. Hence, cloud security is an example of *double* moral hazard.

Externalities come in two varieties: positive and negative. A negative externality occurs when the action of an individual or firm imposes a cost on a third party that does not participate directly in the transaction. Environmental pollution is a classic example [8]. Suppose a factory produces widgets, but as a consequence of producing those widgets, they dump sludge into the river. If the buyer does not have to account for the costs of this pollutant in the price they

pay for the widget, the lower price will stimulate higher demand and create excess pollution and harm to society.

Information insecurity often exhibits negative externalities. Botnets provide the best example [58]. When computers are infected with malware and recruited into a botnet, the harm is not restricted to the operator of the compromised computer. In fact, computers in botnets can be used for a variety of purposes: to send spam, to infect other computers, or to launch denial-of-service attacks. In these cases, the harm is borne by someone other than the owner of the infected computer. Consequently, the incentive for users participating in a botnet to clean up is weakened because they do not fully experience the harm themselves.

The proliferation of data breaches can also be partially explained by negative externalities. In 2017, the US credit bureau Equifax experienced a data breach in which the credit reports of 143 million customers were exposed [10]. While Equifax was responsible for protecting the confidentiality of such sensitive data, they did not bear the full cost of the harm resulting from the disclosure. Many harms affected people and organisations beyond Equifax [38]. As well as the individuals whose data was compromised, other financial institutions and healthcare organisations might have been impacted by increased fraud. Moreover, U.S. government was adversely impacted, primarily through weakened national security resulting from the compromise of sensitive personal financial information of employees with security clearances. When the potential harms of insecurity affect individuals and organisations other than the one who is responsible for taking precautions, we should not be surprised by underinvestment in defensive countermeasures.

A positive externality is a benefit to a third party as a consequence of another's actions. But why is this a market failure? Is an extra benefit to a third party not a good thing? The reason positive externalities causes problems is that the people or firms involved in the transaction do not capture the full benefits and therefore undervalue the transaction. This is harmful because many security investments generate positive externalities by reducing risks for others. For example, if more firms protect themselves successfully against ransomware, there will be fewer successful ransomware gangs. Each firm that successfully protects itself lowers by some small amount the risk to all firms, where the firms who take fewer countermeasures could free ride on the investments of others.

The discovery of software vulnerabilities also creates external benefits. Security researchers work hard to find bugs and disclose them to the software vendor. This benefits the vendor and, downstream, all users of the software. The private benefits of the researcher might be to gain reputation or to learn new skills. The vendor might have a bug bounty program that rewards researchers for finding the bug. Yet in no case will the researcher reap all the downstream benefits. As a result, we should expect less effort devoted to vulnerability discover than is socially optimal.

Information sharing also exhibits positive externalities. Any firm with a security team that monitors cyber attacks against their infrastructure naturally gathers threat intelligence data that might benefit other organisations. Sharing the threat intelligence would create a positive externality.

Moreover, positive network externalities tend to create dominant platforms [33]. A platform is simply a system of interconnected users, such as a telephone network, a social network, or even an operating system. In each case, the network becomes more valuable as it adds more users. We see this with Facebook, where each new person that joins creates some new, albeit minor, additional value to all of the other people who are already on the platform. And

they did not have to do anything to get this benefit. Such positive network externalities are captured in Metcalfe's Law, which asserts that the value of a network grows with the square of its size [37].

When positive network externalities are present, a few dominant platforms tend to emerge that crowd out new entrants. This has serious implications for security, in particular. There is a very significant market reward to pushing out products quickly and ignoring security until the firm establishes itself as the dominant platform. For example, the Windows operating system was quite insecure until well after Microsoft achieved market dominance. Windows XP had lots of significant security holes. With a dominant position secured, Microsoft belatedly invested heavily in security when developing the next version called Vista [45]. Unfortunately for Microsoft, they learned that it is hard to add security in after the fact. Vista was delivered late and over budget [44]. While security improved, significant problems remain to this day. Moreover, dominant platforms exhibit correlated risk. For example, there was a huge spike in ransomware targeting organisations running the same vulnerable Windows file sharing network protocol. Had there been more diversity in the software running on enterprises, such attacks might not be so pervasive today.

Both positive and negative externalities are bad from an economic perspective. With positive externalities, we tend to have less of the good than is desirable. With negative externalities, we end up with more of the bad thing than is desirable. In other words, in a world rife with externalities, we end up with less security investment from the good guys and more harm emanating from the bad guys than would be socially optimal.

The presence of a market failure justifies a regulatory intervention and in turn informs how public policy should be designed. Even when public policy interventions are politically impractical, pointing out the existence of a market failure is still useful for two reasons. First it helps explain why we have suboptimal investment in cyber security. Some puzzles as to why things don't work can be explained in the context of these failures. Second, it could create opportunities and guidance for private actors to come in and correct the problem.

3 MEASUREMENT

A second major thrust of research in security economics over the past two decades has focused on measuring security. Measurement is naturally connected to studying the harms resulting from insecurity. Information asymmetries can arise from a lack of good measurement, such as not documenting the prevalence of attacks or the impact of such attacks, or assessing the effectiveness of countermeasures. We organise this section around two influential frameworks that have systematised the literature to date [5, 6, 60]. These papers contain extensive citations to relevant literature. We encourage readers who might be interested in a deeper dive to check out those papers and associated references.

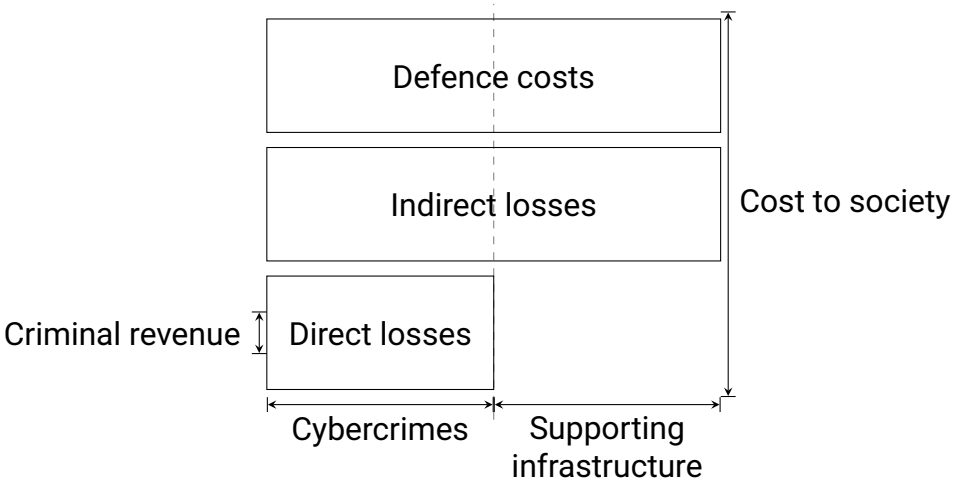

Figure 1: Framework for analysing the costs of cybercrime [5, 6].

3.1 Measuring harms

Harm can take many forms, including physical, psychological, reputational, societal and financial harms [1]. Sometimes these harms can be easily expressed as financial costs, such as lost revenue following an interruption in service. Other times, the harms cannot be easily quantified in financial terms. We refer the interested reader to [1] for an extensive discussion of the various categories of cyber harms. Here, we focus on costs resulting from cybercrime, which is a substantial subset of all cyber harms.

How much does cybercrime cost the economy? It is an easy question to pose, but hard to answer accurately. Many industry estimates have put the cost of cybercrime in the trillions of dollars, yet close inspection usually reveals either questionable or no methodology backing the estimate. It turns out that it is hard to effectively quantify the harms from cyber insecurity. One reason why is that people often conflate different categories of cost, combined with an incentive to hype the problem.

Figure 1 visualises a framework for different components of cybercrime costs. Criminal revenue covers the gross receipts of crime. Direct losses cover losses, damage or other suffering experienced by the victim as a consequence of a crime. By contrast, indirect losses include losses and opportunity costs imposed on society because of certain cybercrimes. The key distinction here is that direct losses affect the victim, while indirect losses cover non-victims. It is a common mistake to equate indirect with the intangible. The reputational damage to a firm that suffers a data breach is a direct loss, even though it is an intangible cost. By contrast, if news of a data breach targeting a retailer makes fewer people shop for fear of their information also being breached, that would be an indirect cost. Defence costs cover the money spent on prevention and controls. The supporting infrastructure to enable multiple types of cybercrime includes botnets, hacked websites and Internet infrastructure operated directly by malicious actors to perpetrate attacks. These are distinguished from the other categories so as to avoid duplicating the costs accounted for. Note that indirect losses span both cybercrimes directly and the supporting infrastructure.

The authors in [5] reported the best estimates in the open literature for the financial costs associated with a wide range of cybercrimes. They find that indirect losses and defence costs typically dwarf direct losses across cybercrime types. Notably, the authors of the study do not report a single financial figure to estimate the total harm from cybercrimes. This is

because individual crime categories have wide ranging confidence intervals, and combining them would yield an unacceptably large range of possibilities. The authors also point out that many important threat types, such as espionage, critical infrastructure attacks and intellectual property theft, lack sufficient data to provide any quantitative estimate of harm at all. Hence, answering the question of what cybercrime costs the economy in total remains an important open question.

3.2 Measuring security effectiveness

We have already seen some reasons why measuring security hard. Firms do not like to discuss openly when they are attacked. Security vendors do not like to subject their tools to independent testing. Even setting aside issues of incentives and market failures, measurement remains a daunting task. One reason why is that security cannot usually be directly measured. Another is that improvements to security reduce the likelihood of being attacked, but other factors are at play too, such as the value of the target to an attacker. Attempts to directly compare security investment with attack likelihood often find a *positive* relationship, i.e., spending more on security is associated with being attacked [48]. Such associations can be spurious if they do not account for the timing of the investment or the value of the target. One study that did control for the timing of security investments and the industries targeted found that firms who adopt more countermeasures do in fact reduce the likelihood of experiencing a cyber incident [23].

Figure 2 presents a causal model that disentangles several factors that impact security effectiveness. Note that there are now two stages to go from threats to harm. First, there is a question whether an attack leads to a successful compromise. Second, a successful compromise can lead to differing levels of harm. At each stage, the security levels of defenders and the exposure of the victim mediate the outcome.

In practice, the security level cannot be observed directly. Instead, a number of observable indicators can be combined to estimate the security level. These indicators might include the adoption of security controls, the presence of open network ports in an enterprise, or hiding webserver version information. It is not the indicators that affect the likelihood of a successful compromise, but rather the unobservable security level as measured by the indicators. For example, Tajalizadehkhoob et al. gathered 15 indicators of security at shared hosting providers and used latent factor analysis to construct indirect measures of security at the hosting provider [52]. Liu et al. gathered publicly-observable network misconfiguration data and compared it to observed malicious activities emanating from that network [35]. Nagle et al. counted open ports to measure a firm's network exposure and linked it to the incidence of botnet activity [40]. These cases illustrate how security level can be indirectly measured from composite indicators.

When examining the literature, two outcome variables are frequently studied. In one, compromise is the outcome of interest. In others, compromise is an input variable, with the resulting harm the outcome variable. The most frequently investigated harm has been the stock price of publicly traded firms that experience data breaches or other compromises. Woods and Böhme identify 16 such event studies [60]. All but one study find weakly negative cumulative abnormal returns, that is, the stock price falls modestly following the public announcement of a breach. More recently, authors have investigated the link between cyber incidents and bond pricing, which can capture longer-term effects as well as apply to a broader population. One study looked at municipal bond prices, finding that IT investment increases after an incident

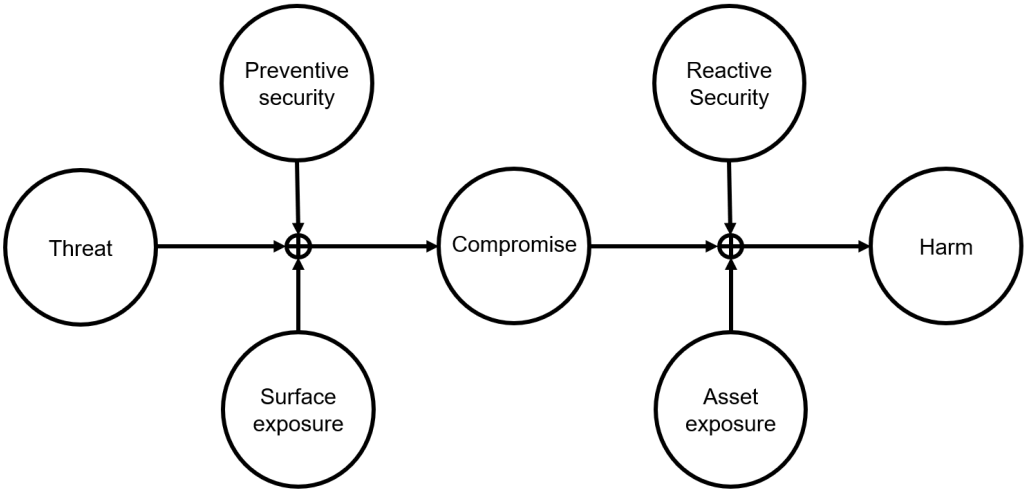

Figure 2: Causal model for measuring security effectiveness [60].

and bonds yields fall two years later [32]. In other words, the municipalities are seen as *less* risky following a ransomware incident. Another study of the broader bond market found that bondholders lose wealth in both the short and long term when firms experience cyber attacks [31].

Apart from changes in stock prices, very few harm studies use financial loss as the outcome variable. Hence, it remains a largely open empirical question as to the extent to which cybersecurity investment by organisations actually impacts the financial harms resulting from cyber attacks.

4 FIRM-LEVEL SOLUTIONS

Firms face many decisions in how to manage cybersecurity in their organisations. Standard risk management approaches have been adapted to cybersecurity, so that cyber risks are treated like any other business risk. The four categories of risk management are acceptance, mitigation, avoidance, and transfer. Most cybersecurity technology can be interpreted as risk mitigation. Cyber insurance is a form of risk transfer. Ideally, firms mange their risks using a combination of all four categories.

Since its inception, security economics researchers have developed models and methods to help firms answer the question of how much security spending is enough. One approach is to develop metrics. A significant challenge with these investment models and metrics is how to quantify the benefit of security. While some may lament that cybersecurity is doomed to be a cost centre, a simple change of perspective can help. Typically, these models define security benefits as the expected loss prevented by having security controls in place. The return on security investment can be calculated as the reduction in expected losses divided by the money spent to achieve the reduction. For these and other metrics, see [11]. See also Section 6.6 of the CyBOK Risk Management and Governance knowledge area for further discussion of security metrics [14].

Another challenge with security metrics is how to obtain reliable data. Typically, metrics require an estimate of losses associated with successful attacks and the probability of loss, both before and after an investment is made. In some cases, researchers construct best-effort point estimates based on available public data. For example, Papa et al. estimated

the losses and probabilities associated with attacks on wastewater facilities using public data [43]. One way to cope with uncertain data is to conduct cost-benefit analysis and identify breakeven points for parameters with unknown values. Another is to construct datasets on losses and probabilities by systematically surveying experts [22, 30]. Nonetheless, it remains an open research problem to collect suitable empirical data to apply security metrics in a robust fashion for many cybersecurity scenarios.

A related effort has been to develop models for security investment. The most influential model of this nature was developed by Gordon and Loeb [25]. The model incorporates a breach probability function that relates a security investment level and inherent vulnerability parameter to the probability of experiencing a loss. The function is chosen so that as spending increases, the marginal benefit in terms of improved security decreases. Such decreasing marginal returns make intuitive sense, particularly when the most cost-effective controls are implemented first. From there, the model identifies the point at which security investment is optimised. It selects the point at which the marginal cost of additional security investment equals the marginal benefit of reduced expected losses.

The Gordon and Loeb model has been influential for several reasons. It established that limitless cybersecurity budgets, even if somehow feasible, are not actually desirable. Instead, there comes a point at which there can be too much security investment. The model produced a rule of thumb, now known as the Gordon-Loeb rule, stating that organisations should never spend more than 37% of their expected loss on security.

It is worth noting that the model, and most of the refinements and extensions that came afterwards (e.g., [26, 27, 50]), are conceptual in nature. They help explain relationships between security investments and outcomes. They are not designed to be directly applied to assist a particular organisation with how to set its exact security budget. The breach probability functions discussed in the paper have been selected for mathematical convenience. It remains an open research question to devise models that have been empirically validated and can be used to guide specific practitioner decisions.

Moreover, it is worth noting that the Gordon-Loeb model adheres to a decision-theoretic perspective, rather than a game-theoretic one. Game theoretic models study strategic interactions between players. In security, those players are typically attackers and defenders. These models can identify equilibrium outcomes when attackers and defenders operate simultaneously, which offers a more realistic perspective on such interactions (see, e.g., [59]). There is a rich literature applying game theory to security that the interested reader is invited to pursue.

If firms are not using security investment models to guide their decision making, what do they do instead? Chief information security officers (CISOs) rely on process-oriented frameworks such as the NIST Cybersecurity Framework, COBIT, or Critical Security Controls [39]. These frameworks sidestep questions about estimating probabilities and expected losses. Instead, they lay out a number of security controls that can be adopted at various levels of sophistication. In essence, these frameworks codify best practice in a structured manner with unified terminology. CISOs who are effective in securing resources often report that security metrics can be helpful when they frame security investments in language chief financial officers (CFOs) understand, like return on investment (ROI) and net present value (NPV).

Cyber insurance has long been advocated not only as a way to manage risk, but also as a private means of solving many of the security failures identified above. Insurers are naturally incentivised to collect data on security investments and their effectiveness to improve under-

writing. Once they have access to such data, they can make evidence based recommendations on which security controls are most effective. Might cyber insurance mitigate the information asymmetries that have plagued the cybersecurity industry? While the potential is undoubtedly there, reality has been less clear cut [61]. Cyber insurance policies grew in prominence only after data breach notification laws came into force, requiring companies to disclose when attacks took place. For other types of cyber attacks not covered by notification requirements, some firms prefer to hide that they have been hacked rather than file an insurance claim, out of concern that this could damage their reputation. It could be hard for insurers to accurately observe the security practices of insureds, particularly during underwriting. The state of the art approach is a detailed questionnaire asking about security practices within the organisation. Unfortunately, these answers are not always accurate, either at the time when answered or as time passes, which weakens the utility of these questionnaires in adjusting premiums to match risks. For a deeper discussion of data quality issues in cyber insurance underwriting, see [42]. Interested readers can also refer to the discussion of insurance in Section 5.4 of the CyBOK Security Operations and Incident Management knowledge area [17]. Nonetheless, the market is maturing, and the potential for cyber insurance to help cybersecurity practice become more empirically grounded remains high.

5 MARKET-LEVEL SOLUTIONS

Thus far we've discussed how to measure security, which helps to quantify these problems and progress towards tackling them. We then discussed different approaches firms can take to manage cybersecurity in their organisations. We now investigate solutions from the perspective of public policy.

This leads to the question: why do we need public policy to solve the security problems identified here? As we have already seen, many cybersecurity issues could be solved independently. However, there are circumstances in which firm-led solutions will fall short. Policy is needed to set the boundaries of unlawful behaviour involving cybersecurity and crime. Policy interventions are needed when market failures like information asymmetries and externalities are present. Governments are useful to solve coordination problems. They are also needed to enforce policies.

Most policy interventions are mandatory. This includes traditional safety regulation and ex post liability. It also includes less traditional approaches, such as certification and information disclosure. Note that there are also policy options available that are voluntary in nature. This includes standardising practices, as well as governments convening and coordinating relevant private actors. Incentives can also be used as part of regulations, such as public procurement rules, to encourage more secure behaviours.

5.1 Ex ante safety regulation

Ex ante safety regulation attempts to stop bad outcomes from happening in the first place. Because regulations impose costs and are never implemented perfectly, the harm we'd like to avoid had better be big, and the cost of cleaning up should the harm materialise must be high. The most extreme example would be nuclear accidents and attacks. These are steps worth taking to prevent failures that would be catastrophic and would have hugely expensive cleanup. In some contexts, the costs of cyber attacks could also be sufficiently high and remedies imperfect and difficult enough that ex ante safety regulation should be considered.

One natural area that fits the bill is critical infrastructure, whose definition seems to call out for ex ante precautions. These are systems and assets whose incapacity or destruction would have a debilitating impact on national security, public health, etc. One might think that this would be the very area in which ex ante cybersecurity regulations already existed. Indeed, most attempts to regulate cybersecurity started with critical infrastructure. Yet in practice we have mostly seen policies that rely on voluntary action and public private partnerships with industry. Why? There are several explanations. First, governments are simply not in a position to articulate the appropriate security requirements. Second, governments do not want to foot the bill for the big investments required, so instead they pursue the more economical approach of pressuring private firms to "do the right thing". Having said that, this is a fast-moving area, and governments across the world have been actively considering legislation that would bring some form of ex ante safety regulation that promotes cybersecurity in critical infrastructures. So watch this space.

A few critical infrastructure sectors, like finance and health, do have some ex ante safety regulations in place for cybersecurity, as these industries have been more heavily regulated for a long time. Where we do see ex ante regulation, it takes a specific form. It is not actually regulating security directly, but effort. Since we cannot measure security like we can measure pollution, regulation cannot require outputs, like mandating a certain level of security. Instead it focuses on inputs, i.e., controls, that we hope correlate with outputs, i.e., security. This is what "maturity levels" are about. And even there, we rarely see inputs in the form of specific technological measures or solutions. Rather, we see procedural norms.

On the positive side, ex ante safety regulation can potentially prevent bad outcomes from happening in the first place, which is important because the costs of cleaning up attacks can often be very high. These regulations also set a baseline of acceptable behaviour that all firms must follow. This can level the playing field as well as achieve better security in weakest link environments [12, 59] and where externalities are present.

But there are downsides too, not least of which is that by setting a minimum standard, there is a risk that firms will be discouraged from doing anything more than the bare minimum. Firms that might do more on security could be inspired to do less if presented with a lax regulation. Moreover, this floor usually sets out procedural measures, rather than security itself. It can be politically challenging to implement ex ante regulations very broadly. Most critical infrastructures do not have any ex ante safety regulations in place, despite these industries fitting the criteria for ex ante regulation by definition. Regulatory capture can also be a problem, when the industry under regulation can set the rules to their own advantage. In cybersecurity, which has a deeply technical aspect, the risks of such capture are especially apparent since rule-setters often lack technical background and seek advice from experts who may push their own agendas. Finally, once regulations are set, adjusting them as technology changes can be hard to pull off.

5.2 Ex post liability

Ex ante safety regulation attempts to change behaviour to prevent a bad outcome from happening. In ex post liability, by contrast, one waits for the bad thing to occur, then assign liability to the party that caused the problem [49]. The idea is that the party anticipates this consequence and hence acts ex ante to prevent the thing from happening – or at least to minimise it to the extent that this is economically rational. Ex post liability has been used extensively in many sectors, such as the auto industry, but it has not yet been widely adopted in software or security. One benefit of ex post liability is that it does not require a regulator to set up rules in advance. Instead, it leaves the decision of how best to avoid harm up to the liable party. In situations where regulators may not know the best course of action, as can happen in specialised areas like cybersecurity, liability offers an elegant solution.

A key question in liability is to whom it should be assigned. Fortunately, economists have a ready answer: the least cost avoider [24]. And just who is that? The least cost avoider is the party to a transaction who incurs the lowest cost to avoid the harm in the first place. It may not be the one who is at fault!

Done right, liability assignment can fix many problems. First, it can often correct misaligned incentives. The party who is made liable for a security failure certainly will be motivated to take steps to prevent the failure from happening. Second, information asymmetries can sometimes be remedied by assigning liability to the party that has better information. Third, externalities arising from insecurity can often be internalised when someone is assigned responsibility for the problem.

Let's now illustrate liability assignment by talking about software. Software has always had bugs, and many high-profile security incidents can be traced back to vulnerabilities in software, such as Eternal Blue [13], the 2017 Equifax data breach [10] and 2020 SolarWinds hack [47]. It is undisputed that many software developers do not follow secure development practices or take enough steps to detect and remove vulnerabilities before shipping code. Software liability would place the burden of insecure code directly on the developers. To proponents, this makes sense because software developers are in the best position to fix the vulnerabilities at the lowest cost — they are the least cost avoiders. Eliminating a vulnerability at its source is cheaper and more effective than expecting the thousands or millions of customers running the code to take safety precautions to block the impact of attacks exploiting the vulnerability. Assigning liability to developers would undoubtedly incentivise investment in secure software development practices, which has been lagging for decades. Memory-safe languages have been available for decades, yet production-level software continues to be written in C, and vulnerabilities exploiting memory bounds-checking failures continue to be discovered and exploited in new code.

Meanwhile, software liability has considerable downsides. For a start, writing bug-free code today remains impossible. Software engineering practices have not matured to the level where it is reasonable to expect even proficient programmers to introduce zero vulnerabilities into their code. Opponents of software liability argue that it is therefore unfair to blame developers when vulnerabilities are found. A related argument is that software liability could hinder innovation, especially in the context of developing free and open source software, and that much open source software development would stop if those volunteering their efforts could be held liable for bugs. This argument is not convincing, since such software is already exposed to liability when used in products covered by liability rules. Opponents also argue that software liability would raise the barriers of entry to software development, meaning that

small firms may not enter the marketplace.

To sum up, ex post liability is a powerful tool to align incentives, potentially correcting market failures along the way. It can be a good option if the costs of cleanup are manageable and if liability can be clearly assigned to the least cost avoider. Liability is not without controversy, due in no small part to its potency. Technology providers have managed to disclaim liability for the safety of their products by arguing that liability would have a chilling effect on innovation. It remains to be seen if such arguments will continue to pass muster as society deepens its technological dependence.

For more discussion of liability in cybersecurity, see the CyBOK Law and Regulation knowledge area [15].

5.3 Certifying products

A number of mechanisms can be used to remedy information asymmetries. One approach is to use certification schemes, which evaluate products independently to establish their security. While it may be hard for end users or customers to evaluate security, the information asymmetry could be remedied by turning the problem over to an expert body to pass judgement. We see product certifications in many contexts from Underwriters Laboratory to the ubiquitous CE mark, which is intended to demonstrate that the labelled product has been self-certified by its manufacturer to meet EU health and safety requirements [21]. Could similar schemes work for cybersecurity?

This is not a new idea. The Orange Book originated in the 1970s and 80s as a means of certifying the security of products developed for the US Department of Defense. After related efforts by Britain and Germany, the Common Criteria was established in the 1990s as an international standard for certifying computer security. The system does work in that products that go through this process do have a modicum of security baked in, but care is needed in the design and implementation of the certification process. There can be issues with the incentives for the evaluators. In some cases, the evaluation can be paid for by the vendor seeking approval, and so vendors can shop around for the evaluator willing to give the easiest ride on the testing.

That's effectively what happened in a case involving PIN entry devices in the United Kingdom that had been Common Criteria evaluated to be tamper-evident. Despite successful certification, researchers were able to insert a paper clip to tap the board and read the PIN as it was sent over the wire, without being detected [18]. The evaluation completely failed, even though it was evaluated to be secure, in practice it was shown not to be.

Research on website security seals has found that criminals are attracted to set up these seals in order to convey a false sense of trust. Ben Edelman has looked at data from SiteAdvisor to find lists of bad websites distributing malware. He then looked at websites that signed up for TrustE seals, finding that bad websites were more than twice as likely to be TrustE-certified than good websites [19]. Criminals are co-opting the seal, which means we end up in adverse selection where the bad participants are more likely to seek the seals than good – this is the exact opposite of the intended goal of the certification schemes.

The appeal of certifying products to be secure is not going away. Recently, governments across the world have started pushing for cybersecurity labelling of devices related to Internet of Things (IoT) devices. The Mirai botnet was assembled with over 100,000 hacked security cameras and other embedded devices [7]. Today, consumers can compare these devices

on price and features, but not security. Labels could potentially change that. Singapore was the first country to launch an IoT cybersecurity labeling program [16]. As should be clear by now, the devil is in the details. Ensuring that the labels accurately measure cybersecurity and communicate it clearly to consumers is essential, and likely to be a work in progress in the years ahead.

5.4 Certifying processes

Why certify processes? For the same reasons one might certify a product. There is often an information asymmetry that exists between organisations that interact. When these interactions involve the sharing of personal data or reliance on products and services, it is reasonable to worry about the security of the other firms' operational practices. Supply chain security depends critically on the security of third-party suppliers, and often these suppliers provide services rather than products. The key question is how to be assured that the business processes employed by outsiders are in fact secure.

ISO 27001 is an international standard for managing cybersecurity. Organisations can apply for certification and go through an assessment and audit process to verify that baseline standards are being met. There is nothing particularly earth-shattering in the ISO requirements, but they do provide a base level for secure practices. Sometimes contracts require vendors to be ISO-certified; indeed, obtaining an ISO 27001 certification can provide a meaningful signal to prospective customers and partners that cybersecurity is taken seriously by the organisation. Can ISO-27001 certified companies still be hacked? Of course, because the certification requirements are not particularly strenuous. It is in the interest of the standards body to make the standards achievable, because certifications derive much of their value from their ubiquity. More certifications means more fees for the certification issuer, growing influence as a signalling device, and the potential for declining quality . Moreover, what is being certified is that companies comply with expected inputs to security, such as implementing certain controls, rather than ensuring the companies themselves are secure, as the latter cannot be directly measured.

5.5 Information disclosure

Let's discuss another strategy for correcting information asymmetries: information disclosure. We start with a powerful example of information disclosure that directly inspired similar efforts in cybersecurity. Like cyber attacks, environmental pollution can be difficult to observe when it happens. Addressing this was the idea behind the Toxic Release Inventory (TRI) administered by the US Environmental Protection Agency (EPA) [53]. Firms regularly dump pollutants, whether intentionally or by accident. Often, it is not even illegal when the spill happens. The law supporting the TRI created a requirement to disclose when, where, and how much pollutants are spilled. The EPA then aggregates the reports and shares them with the public, so that people become aware of any risks where they live and can take action if necessary. Note that the law did not introduce fines or penalties for pollution. It simply made disclosure mandatory and shared the information with the public. What happened? Pollution rates fell! The mere threat of embarrassment was enough to change behaviour. The TRI had a second effect. Communities became aware of when and where pollution was happening, which empowered them to take appropriate precautions and to advocate for change.

Let's return to cybersecurity. Back in 2002, California passed legislation requiring companies operating in the state to disclose any breach of security that left a California citizen's personal

information exposed [51]. This clause was inserted into a larger bill at the behest of Deirdre Mulligan, who was working with the state assembly at the time, and is now a professor of Information at the University of California Berkeley. Prof. Mulligan was inspired by the TRI, arguing that consumers should be similarly empowered to take precautions when their personal information was breached. She also argued that companies might similarly reduce the amount of personal data pollution dumped online. She convinced a legislator to insert the clause into the law and it stuck. As a result, the past decade has seen a cascade of data breach reports which had arguably been going on for years prior to the disclosure requirement. Now these laws are on the books or under consideration in many more US states and countries, including the European Union.

What has been the effect of data breach legislation? Before the legislation, data breaches were not taken seriously as a risk by many companies. Now, it has the full attention of boards [9]. This illustrates a more general principle: many hard security problems can actually be managed if we can only measure the risks appropriately and assign responsibility when things go wrong. Before the legislation, data breaches didn't meet these criteria, and so it seemed an intractable problem. Now, the scope of the problem has been made abundantly clear, companies will sell insurance policies, and firms know what they need to do in order to protect personal information they hold.

What other aspects of cybersecurity might benefit from greater information disclosure? Financial fraud figures are often not reported, which makes it hard to estimate the true magnitude of cybercrime risks. Cyber espionage is fiendishly difficult to detect, and when it is discovered, firms usually prefer not to admit it publicly. Control systems that manage critical infrastructure may be threatened by hackers, yet we don't really know the true extent because reports have been sporadic and non-attributed. More broadly, a more comprehensive and ongoing collection of data on cybercrime losses would help society rate its improvement in managing cybersecurity overall.

Another example of information disclosure comes from the US Securities and Exchange Commission (SEC), which issued guidance that it expects publicly traded firms to disclose all material cyber incidents in their regulatory filings [55, 56, 57]. This is significant, because it covers incidents of any type, not just data breaches. The SEC guidance illustrates another general lesson about information disclosure: without standards on the shape disclosures should take, the resulting information will be messy and hard for others to learn from. Firms are free to discuss breaches in their regulatory filings as they see fit. These documents are typically verbose, full of legalese, and written in such a way to reflect as favourably as possible on the company. It can be hard to identify when attacks are being reported and to determine the significance [29]. The contrast with the TRI, which is reported to the EPA in a standardised format and shared in a public database, is striking.

The final information disclosure effort we'll discuss is very different than the ones presented so far. A software bill of materials (SBOM) provides a machine-readable record of external packages and libraries utilised in software [41]. Without an SBOM, it can be impossible to identify vulnerable dependencies, particularly for proprietary software. Hence, SBOMs can help correct the information asymmetry that exists in software when it comes to observing its security.

Compared to other policy interventions such as safety regulation and liability, information disclosure is not very ambitious. It simply compels disclosure of relevant information, such as security incidents, but does not assign further blame or require any change in behaviour. Any security problems involving externalities are not helped. Even so, it is widely believed that

underreporting of cybersecurity remains a significant problem, so we may not be getting the full picture. Despite these limitations, it is a good start. information asymmetry is a critical market failure affecting cybersecurity, and information disclosure tackles it head on. Data breach notification requirements have brought cybersecurity to the attention of the C-suite, and have prompted a substantial investment to mitigate the risk. We have seen that this light-touch policy intervention can foster significant private sector investment and change business practices.

5.6 Government failures

Governments play a key role in improving cybersecurity, primarily by remedying market failures using the policy interventions discussed above. Nonetheless, there is a risk of government failures, which can arise when a government intervention causes a more inefficient allocation of goods and resources than would occur without that intervention. Precisely when governments are trying to improve security in the face of market failures, they might end up exacerbating the problem or create new ones. We briefly discuss three canonical government failures that can affect cybersecurity: regulatory capture, rent-seeking and regulatory arbitrage,

Regulatory capture happens when an agency advances the commercial or political interests of the group that it is supposed to regulate. Regulatory capture naturally arises in sectors where firms possess expertise on industry operations that outsiders lack. When firms provide that expertise, there is a substantial risk that the resulting rules will favour them. We can see the risk of regulatory capture through the many lobbyists employed by big tech firms. Regulatory capture is not unique to tech firms. Consider that taxi regulations in many E.U. countries entrenched the existing industry, which raised significant barriers to new entrants like Uber and Lyft. In general, we often see demands for expensive compliance regimes and regulations being favoured by established firms who may have captured regulators. It is why dominant firms have called for regulation of social networks and artificial intelligence. It is also helps explain why large firms tend to be more supportive of cybersecurity regulations.

Rent-seeking behaviour seeks to extract some of the legitimate value generated by economic transactions, often through regulation. A classic example of rent-seeking is the tendency of the security industry to inflate the severity of threats. Another example is the frequent suggestion to promote audits or certifications to promote security. Large consulting firms might recommend ISO 27001 certification, or similar schemes, while at the same time selling services to help implement the certifications and audit implementation.

Regulatory arbitrage exploits differences in regulations in two or more regions to bypass unfavourable regulation. For example, many global tech companies select Ireland for its E.U. headquarters, as it is widely viewed as having the most “business-friendly” privacy enforcement regime. Cryptocurrency exchanges regularly set up operations outside the U.S. in order to offer riskier financial products like derivatives and highly leveraged trading in a bid to avoid U.S. financial regulators. But legislators have fought back, as evidenced by the E.U.’s General Data Protection Regulation which covers any company that processes the personal data of EU subjects, regardless of where the company is based [20].

Issues involving regulation are discussed at great length in the CyBOK Law and Regulation knowledge area [15].

6 CONCLUDING REMARKS

As cybersecurity has grown in importance in recent decades, scholars have realised that a wide range of disciplinary perspectives beyond computer science are valuable. Social sciences in general, and economics in particular, provide a useful lens to describe many of the fundamental challenges as well as identify sound approaches to tackle them. This knowledge guide has shared some of the key insights made possible through economic analysis.

At the most basic level, incentives matter, both for attackers and defenders. When a system is not protected or an unthinkable attack is reported, look at the incentives first. While there undoubtedly will be technical explanations for failure, the root cause is more likely to be economic in nature. Cyber insecurity can also often be explained by the presence of market failures, notably information asymmetries and externalities. Recognising and naming these failures is an important first step. Next is to seek countermeasures that can correct market failures. Mandatory breach notification laws and liability to internalise externalities are two promising approaches.

Looking ahead, economics is bound to continue to play an important role in improving cybersecurity writ large. By blending technical and economic approaches, the state of cybersecurity may ultimately be improved.

REFERENCES

[1] Ioannis Agrafiotis, Jason R. C. Nurse, Michael Goldsmith, Sadie Creese, and David Upton. A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. *Journal of Cybersecurity*, 4(1):tyy006, 10 2018.

[2] George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

[3] Ross Anderson. Why cryptosystems fail. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, pages 215 – 227, New York, NY, USA, 1993. Association for Computing Machinery.

[4] Ross Anderson. Why information security is hard - an economic perspective. In *Seventeenth Annual Computer Security Applications Conference*, pages 358–365, 2001.

[5] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Carlos Gañán, Tom Grasso, Michael Levi, Tyler Moore, Stefan Savage, and Marie Vasek. Measuring the changing cost of cybercrime. In *18th Workshop on the Economics of Information Security (WEIS)*, 2019.

[6] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michael van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In *11th Workshop on the Economics of Information Security (WEIS)*, 2012.

[7] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. Understanding the mirai botnet. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1093–1110, Vancouver, BC, August 2017. USENIX Association.

[8] Robert U. Ayres and Allen V. Kneese. Production, consumption, and externalities. *The American Economic Review*, 59(3):282–297, 1969.

[9] Kenneth A. Bamberger and Deirdre K. Mulligan. Privacy on the books and on the ground. *Stanford Law Review*, 63(2):247–315, 2011.

[10] Tara Siegel Bernard, Tiffany Hsu, Nicole Perlroth, and Ron Lieber. Equifax says cyberattack may have affected 143 million in the U.S. *New York Times*, September 2017. **https://www.nytimes.com/2017/09/07/business/equifax-cyberattack.html**.

[11] Rainer Böhme. Security metrics and security investment models. In Isao Echizen, Noboru Kunihiro, and Ryoichi Sasaki, editors, *Advances in Information and Computer Security*, pages 10–24, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[12] Rainer Böhme and Tyler Moore. The iterated weakest link. *IEEE Security & Privacy*, 8(1):53–55, 2010.

[13] Thomas Brewster. An NSA cyber weapon might be behind a massive global ransomware outbreak. *Forbes*, May 2017. **https://www.forbes.com/sites/thomasbrewster/2017/05/12/nsa-exploit-used-by-wannacry-ransomware-in-global-explosion/**.

[14] Pete Burnap. *The Cyber Security Body of Knowledge v1.1.0, 2021*, chapter Risk Management & Governance. University of Bristol, 2021. KA Version 1.1.1.

[15] Robert Carolina. *The Cyber Security Body of Knowledge v1.1.0, 2021*, chapter Law & Regulation. University of Bristol, 2021. KA Version 1.0.2.

[16] CSA Singapore. Cybersecurity labelling scheme, 2021. **https://www.csa.gov.sg/our-programmes/certification-and-labelling-schemes/cybersecurity-labelling-scheme**.

[17] Hervé Debar. *The Cyber Security Body of Knowledge v1.1.0, 2021*, chapter Security Operations & Incident Management. University of Bristol, 2021. KA Version 1.0.2.

[18] Saar Drimer, Steven J. Murdoch, and Ross Anderson. Thinking inside the box: System-level failures of tamper proofing. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 281–295, 2008.

[19] Benjamin Edelman. Adverse selection in online "trust" certifications and search results. *Electronic Commerce Research and Applications*, 10(1):17–25, 2011. Special Section: Service Innovation in E-Commerce.

[20] European Commission. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016.

[21] European Commission. CE marking – general guidelines, 2021. **https://single-market-economy.ec.europa.eu/single-market/ce-marking_en**.

[22] Jack Freund and Jack Jones. *Measuring and managing information risk: a FAIR approach*. Butterworth-Heinemann, 2014.

[23] Neil Gandal, Tyler Moore, Michael Riordan, and Noa Barnir. Empirically evaluating the effect of security precautions on cyber incidents. *Computers & Security*, 133:103380, October 2023.

[24] Stephen G. Gilles. Negligence, strict liability, and the cheapest cost-avoider. *Virginia Law Review*, 78(6):1291–1375, 1992.

[25] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Trans. Inf. Syst. Secur.*, 5(4):438–457, nov 2002.

[26] Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. Investing in cybersecurity: Insights from the Gordon-Loeb model. *Journal of Information Security*, 7(02):49, 2016.

[27] Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. Integrating cost–benefit analysis into the NIST Cybersecurity Framework via the Gordon–Loeb model. *Journal of Cybersecurity*, 6(1):tyaa005, 2020.

[28] Andrew Granato and Andy Polacek. The growth and challenges of cyber insurance. *Chicago Fed Letter*, 426:1–6, 2019.

[29] Abulfaz Hajizada and Tyler Moore. On gaps in enterprise cyber attack reporting. In *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages

227–231, Los Alamitos, CA, USA, jul 2023. IEEE Computer Society.

[30] Douglas W. Hubbard and Richard Seiersen. *How to Measure Anything in Cybersecurity Risk*. Wiley, 2016.

[31] Subramanian R. Iyer, Betty J. Simkins, and Heng Wang. Cyberattacks and impact on bond valuation. *Finance Research Letters*, 33:101215, 2020.

[32] Jonathan Jensen and Fiona Paine. Municipal cyber risk. In *22nd Workshop on the Economics of Information Security (WEIS)*, 2023. **https://weis2023.econinfosec.org/wp-content/uploads/sites/11/2023/06/weis23-jensen.pdf**.

[33] Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, 1985.

[34] Eireann P. Leverett. Quantitatively assessing and visualising industrial system attack surfaces. Technical report, University of Cambridge Computer Laboratory, 2011. **https://www.cl.cam.ac.uk/~fms27/papers/2011-Leverett-industrial.pdf**.

[35] Yang Liu, Armin Sarabi, Jing Zhang, Parinaz Naghizadeh, Manish Karir, Michael Bailey, and Mingyan Liu. Cloudy with a chance of breach: Forecasting cyber security incidents. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 1009–1024, Washington, D.C., 2015. USENIX Association.

[36] Andrew Martin, Awais Rashid, Howard Chivers, George Danezis, Steve Schneider, and Emil Lupu. *The Cyber Security Body of Knowledge v1.1.0, 2021*, chapter Introduction to CyBOK. University of Bristol, 2021. Version 1.1.0.

[37] Bob Metcalfe. Metcalfe's Law after 40 Years of Ethernet. *Computer*, 46(12):26–31, 2013.

[38] Tyler Moore. Prepared Testimony of Tyler Moore before the U.S. Senate Committee of the Judiciary Subcommittee on Privacy, Technology and the Law, October 2017. **https://www.judiciary.senate.gov/imo/media/doc/10-04-17%20Moore%20Testimony1.pdf**.

[39] Tyler Moore, Scott Dynes, and Frederick Chang. Identifying how firms manage cybersecurity investment. In *15th Workshop on the Economics of Information Security (WEIS)*, 2016.

[40] Frank Nagle, Sam Ransbotham, and George Westerman. The effects of security management on security events. In *16th Workshop on the Economics of Information Security*, 2017.

[41] National Telecommunications and Information Administration (NTIA). Minimum elements for a software bill of materials (SBOM), 2021. **https://www.ntia.doc.gov/report/2021/minimum-elements-software-bill-materials-sbom**.

[42] Jason R.C. Nurse, Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. The data that drives cyber insurance: A study into the underwriting and claims processes. In *2020 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, pages 1–8, 2020.

[43] Steve Papa, William Casper, and Tyler Moore. Securing wastewater facilities from accidental and intentional harm: a cost-benefit analysis. *International Journal of Critical Infrastructure Protection*, 6(2):96–106, 2013.

[44] Rory Reid. Microsoft Vista late – but why? *CNET*, Mar 2006. **https://www.cnet.com/tech/computing/microsoft-vista-late-but-why/**.

[45] Aaron Ricadela. Gates says security is job one for Vista. *Information Week*, feb 2006. **https://www.informationweek.com/software-services/gates-says-security-is-job-one-for-vista**.

[46] Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.

[47] David Sanger. Russian hackers broke into federal agencies, U.S. officials suspect. *New York Times*, December 2020. **https://www.nytimes.com/2020/12/13/us/politics/**

russian-hackers-us-government-treasury-commerce.html.
[48] Ravi Sen and Sharad Borle. Estimating the contextual risk of data breach: An empirical approach. *Journal of Management Information Systems*, 32(2):314–341, 2015.
[49] Steven Shavell. Liability for harm versus regulation of safety. *The Journal of Legal Studies*, 13(2):357–374, 1984.
[50] Henry R.K. Skeoch. Expanding the Gordon-Loeb model to cyber-insurance. *Computers & Security*, 112:102533, 2022.
[51] State of California. Senate Bill No. 1386, 2002. http://www.leginfo.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.pdf.
[52] Samaneh Tajalizadehkhoob, Tom van Goethem, Maciej Korczyński, Arman Noroozian, Rainer Böhme, Tyler Moore, Wouter Joosen, and Michel van Eeten. Herding vulnerable cats: A statistical approach to disentangle joint responsibility for web security in shared hosting. In *ACM SIGSAC Conference on Computer and Communications Security*, CCS '17. ACM, 2017.
[53] United States. Emergency Planning and Community Right-to-Know Act (EPCRA) Section 313, 1986.
[54] United States. Regulation E: Electronic Fund Transfers 12 CFR 205, 1996. https://www.ecfr.gov/current/title-12/chapter-II/subchapter-A/part-205.
[55] U.S. Securities and Exchange Commission. Cf disclosure guidance: Topic no. 2, 2011. https://www.sec.gov/rules/other/2018/34-83723.pdf.
[56] U.S. Securities and Exchange Commission. Commission statement and guidance on public company cybersecurity disclosures, 2018. https://www.sec.gov/rules/interp/2018/33-10459.pdf.
[57] U.S. Securities and Exchange Commission. Cybersecurity risk management, strategy, governance, and incident disclosure, 2022. https://www.sec.gov/rules/proposed/2022/33-11038.pdf.
[58] Michel Van Eeten, Johannes M Bauer, Hadi Asghari, Shirin Tabatabaie, and David Rand. The role of internet service providers in botnet mitigation an empirical analysis based on spam data. In *Telecommunications Policy Research Conference (TPRC)*, 2010.
[59] Hal R. Varian. System reliability and free riding. In *Economics of information security*, pages 1–15. Springer, 2004.
[60] Daniel W. Woods and Rainer Böhme. SoK: Quantifying cyber risk. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 211–228, 2021.
[61] Daniel W. Woods and Tyler Moore. Does insurance have a future in governing cybersecurity? *IEEE Security & Privacy*, 18(1), 2020.

GLOSSARY

adverse selection consequence of information asymmetries in which the party with more information benefits more than the party with less information.

criminal revenue gross receipts from crime.

defence cost money spent on prevention and controls.

direct loss loss, damage or other suffering experienced by the victim as a consequence of crime.

ex ante safety regulation policy intervention in which risky behaviour is regulated in order to make its occurrence less likely.

ex post liability policy intervention in which responsibility for failure is assigned to a decision maker in order to deter the bad outcome from happening.

externality market failure that occurs when the action of an individual or firm imposes a cost or benefit on a third party that does not participate directly in the transaction.

indirect loss loss or opportunity cost imposed on society because of cybercrimes.

information asymmetry market failure that occurs when when one party in a transaction knows more than the other.

information disclosure policy intervention in which private information must be revealed to mitigate an information asymmetry.

ISO 27001 international standard for managing cybersecurity by organisations.

market failure situations in which reality differs from models of perfect competition, usually one of the following categories: monopoly, oligopoly, public goods, information asymmetries, and externalities.

Metcalfe's Law assertion that the value of a network grows with the square of its size.

monopoly market failure that occurs when a good or service has only one provider.

moral hazard consequence of information asymmetries in which decision makers adopt a riskier behaviour in response to receiving protection from the behaviour's consequences.

oligopoly market failure that occurs when a good or service has only a small number of providers.

public good goods that are both nonexcludable, i.e., there is no practical way to prevent people who don't pay from consuming the good, and nonrival, i.e., when someone consumes the good, others can still consume it too plural.

ACRONYMS

ATM Automated Teller Machine.

CFO Chief Financial Officer.

CISO Chief Information Security Officer.

EPA Environmental Protection Agency.

NPV Net Present Value.

ROI Return on Investment.

SEC Securities and Exchange Commission.

TRI Toxic Release Inventory.