Blessed Are The Lawyers, For They Shall Inherit Cybersecurity

Daniel W. Woods
Leopold-Franzens-Universität Innsbruck
Austria
daniel.woods@uibk.ac.at

Aaron Ceross
University of Oxford
United Kingdom
aaron.ceross@cs.ox.ac.uk

Abstract

This paper considers which types of evidence guide cybersecurity decisions. We argue that the "InfoSec belongs to the quants" paradigm will not be realised despite its normative appeal. In terms of progress to date, we find few empirical results that can guide risk mitigation decisions. We suggest the knowledge base about quantitative cybersecurity is continually eroded by increasing complexity, technological flux, and strategic adversaries. Given these secular forces will not abate any time soon, we argue that legal reasoning will increasingly influence cybersecurity decisions relative to technical and quantitative reasoning. The law as a system of social control bristles with ambiguity and so legal mechanisms exist to resolve uncertainties over time. Actors with greater claims to authority over this knowledge base, predominantly lawyers, will accrue decision making power within organisations. We speculate about the downstream impacts of *lawyers inheriting cybersecurity*, and also sketch the limits of the paradigm's explanatory power.

Keywords: risk management, cyber policy, quantitative cybersecurity, philosophy of science, lawyers,

ACM Reference Format:
Daniel W. Woods and Aaron Ceross. 2021. Blessed Are The Lawyers, For They Shall Inherit Cybersecurity. In *NSPW'21: New Security Paradigms Workshop.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3498891.3501257

1 Introduction

In considering Dan Geer et al.'s argument about why information security belongs to the quants [41], it is important to recall the fact–value distinction. As a normative claim, the authors [41] taps into widespread belief that security decisions should be guided by quantitative evidence. As a predictive claim, it imagines a world with a knowledge base very different to what has been produced by the scientific community thus far [108, 120]. The moralistic fallacy occurs when the persuasive force behind the normative claim about how the world should be spills over into the descriptive/predictive claim about how the world is/will be.

This paper argues that the normative force behind the *security belongs to the quants* perspective has distracted from the immature reality of quantitative cybersecurity. Multiple research surveys fail to identify interventions that influence real-world cybersecurity outcomes [73, 108, 120], and there is little to suggest we are on the cusp of a scientific breakthrough. However, we do not believe the status quo in which decisions are guided by the received wisdom of InfoSec will hold. We put forward an alternative proposition that *lawyers will inherit cybersecurity*, which we justify by reasoning about how knowledge is generated by each field.

Efforts to build knowledge by techies and quants are undermined by factors like complexity, technological flux, and shifting adversaries. Both fields can tolerate the resulting uncertainty—quants accept null results and that some decisions have no corresponding evidence, while security professionals qualify recommendations with statements like 'nothing is 100% secure'. In contrast, law as a system of social control can only tolerate so much uncertainty. There is a secular trend towards resolving ambiguity as regulators issue guidance, case law evolves, and lawyers establish common practice while advising firms.

Over time, knowledge about how to mitigate legal risk will grow more quickly and remain relevant for longer relative to knowledge about mitigating technical risk. This becomes relevant because organisations exposed to both technical and legal risk are likely to prioritise interventions according to relative certainty, as well as the actual effectiveness. Thus, the relative ease of reasoning about legal risk will shift resources away from technical measures. In this sense, techies, quants and lawyers are engaged in a zero-sum battle over who controls cybersecurity decisions. Our descriptive/predictive claim is that power is shifting towards lawyers and will continue to do so—we avoid the normative claim that this shift is an improvement, and tentatively argue the converse.

Having established that lawyers and not quants will inherit cybersecurity, we begin to sketch where our theory is

NSPW '21, October 26–28, 2021, Virtual Conference
© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8573-2...$15.00
https://doi.org/10.1145/3498891.3501257

most applicable and where it breaks down. Certain aspects of cybersecurity are already vulnerable to legalisation like incident response due to the thicket of reporting obligations, whereas others will remain the domain of techies. We also add nuance to our theory of how legal knowledge is generated as this too has imperfections and even indeterminacy conditions [18]. Lawyers inheriting cybersecurity cannot sideline the technical community entirely, rather the legal system will establish new hierarchies of expertise.

The paper is organised as follows. Section 2 identifies normative and descriptive papers related to quantitative cybersecurity with the goal of undermining the proposition that *information security belongs to the quants*. Section 3 sketches how knowledge is generated by techies, quants, and lawyers. We use this to reason about why lawyers can offer comparatively more certainty about cyber risk. Taking the predictive claim as a given, Section 4 speculates about the impact of lawyers inheriting cybersecurity. We discuss the limits and nuances of the theory in Section 5. Section 6 concludes the paper.

2 Does InfoSec Belong to the Quants?

This section outlines the "information security belongs to the quants" as a normative claim (Section 2.1) and then as a descriptive claim (Section 2.2).

2.1 Normative

We can distil normative quantitative security papers into explicit and implicit arguments. Explicit arguments can often be described as *position pieces* and follow the general theme of arguing "we should measure security". Variations on this theme include adding specificity to the domain ("we should measure cloud security") or specificity to the type of measurements ("we should measure security in X way"). Implicit arguments about the normative value of quantitative security are common in theoretical works. For example, representing security decisions via continuously differential production and utility functions (as is common in Game Theory [72]) implicitly assumes decisions can and should be optimised, such as in the most widely cited publication on security investments [44].

Turning to the explicit arguments, Dan Geer et al. [41] invoke Bernstein's history of risk management [9] to contrast two options for InfoSec: (i) the status quo "of oracles and soothsayers" retaining control; (ii) risk management as an objective science guiding decisions. Baker et al. [7] instead frame the choice between uncertainty and quantitative methods. Despite the simplicity of such choices, quantitative methods have not yet taken off as evidenced by a recent research agenda [33].

A common approach is to suggest a promising new approach to measuring cyber risk. For example, Geer et al. [41] find inspiration in finance, such as portfolio management or insurance. Alternative ways forward include Value at Risk models [53], principled data collection [7], Elo-style rating [79], reliability approaches [65], input-output models [54], and pragmatism [50]. The quantitative cybersecurity paradigm can continue generating such normative papers until no new approaches remain.

2.2 Descriptive

An alternative way to evaluate the outlook for quantitative cybersecurity is to survey progress thus far. The following section surveys empirical cyber risk papers. We seek to identify reduced form results about effectiveness of different interventions. Such results can be entirely independent of any underlying theory about how or why the mechanism is effective. For example, one can appreciate that a vaccine with an effectiveness of 95% is preferable to a vaccine with 56% without any understanding of molecular biology.

We classify studies into technical, financial, and legal based on the outcome variable in the study. For example, a technical study might try to explain which devices are infected, whereas financial studies quantify the impact of an incident in terms of stock market value or dollar loss. Legal studies might concern which incidents are litigated using the methods of quantitative social science. Notably, studies in this section do not use legal reasoning. We select from studies identified by prior surveys to illustrate our points. Readers looking for a comprehensive survey should refer to the systematic reviews of each field (computer security [108, 120] and finance [73]).

Technical Studies. Here we unpack the insights from a survey of 30 years of empirical cyber risk research [120]. There are few results speaking to the efficacy of individual security controls in mitigating adverse technical outcomes. In fact, such results often get causal direction backwards due to confounding variables [120, p. 11]. For example, greater security budgets are associated with higher frequency of data breach [92] and web servers with more up-to-date software are more likely to be compromised [106], although applying security updates does lead to better outcomes in terms of *re*-compromise rates.

Edwards et al. [29] provide an illustrative exception in their study. They find that organisations who do not block peer-to-peer file sharing have "318 times higher" rates of botnet compromise, and also find TLS configuration errors are statistically significant correlates of compromise. Finding such results is undoubtedly positive, but it is difficult to see how future work can improve on this infrastructure given the data was collected by an industry-leading proprietary scanning system. Even with this infrastructure, the resulting recommendations—block file sharing and configure TLS correctly—would underwhelm most security managers, especially given the relationship could be correlation rather than causation.

Rather than directly regressing controls on risk outcomes, an alternative approach is to introduce an intermediate variable combining many security controls. A number of papers adopt this approach [66, 101, 124] (though they do so with different mathematical representations). Collectively, the results show that aggregating many indicators of security provides more explanatory power [120]. This comes at the cost of obscuring the causal role of individual controls. The opaque relationship between individual controls and risk outcomes will likely worsen as the ML convention of reporting prediction rates becomes the norm (as opposed to reporting the effect size of explanatory variables in regressions). It is difficult to make recommendations based on such studies.

Finance studies. A 2021 survey of cyber risk management in finance identified a research gap regarding "the effectiveness of countermeasures" [73]. The paper's section on cyber risk mitigation identified: conceptual frameworks [33, 95], a theoretical model [44], and a survey of economic papers [6], but found no empirical studies. This survey serves as a benchmark expectations for finance studies investigating the effectiveness of interventions. We identify three types of study: incident repositories, surveys, and stock market reaction.

Incident repositories collect historic information about the causes and impact of cyber incidents. These databases are populated by: trawling the web for public information as done by Advisen [1, 77, 85], getting data breach information from governments and regulators like Privacy Rights Clearing House do [28, 30, 113] or by receiving private report from members [11, 15, 31]. Although this creates rich information about how cyber losses have varied over time, it is difficult to link losses to technical procedures. Information about the historic security procedures at each organisation who suffered a loss is largely lost to time unless combined with some other data source, which is rarely done. Aldasoro et al. [1] provide an exception by collecting information about security spending and staffing at the sector level, but this exception actually supports the more general point that studies using incident repositories do not quantify the efficacy of preventative security [120].

The second type of study avoids this problem because survey instruments can ask respondents about both losses suffered *and* the defences in place. Biancotti [10] takes advantage of the Bank of Italy adding cybersecurity questions to their annual survey. She warns that the data is not well suited to estimating "how effective defensive expenditure [is]" [10, p. 28], and then shows the preliminary result that defensive expenditure is positively correlated with the probability of suffering a breach. This *more security, more compromise* relationship likely results from a failure to include confounding variables [120]. The outlook for survey research is promising if they solve reporting biases [35], and pose the right questions about security measures in a way that can be independently validated.

The third type of study quantifies how stock market price changes following the announcement of cybersecurity news like a data breach or an investment. The academic literature suggests financial markets reward post-breach informational interventions rather than technical measures. For example, Amir identifies a bias in the reporting of events in which "managers disclose less severe attacks and withhold information from investors on attacks that cause greater damage" [3]. Gay [39] shows that releasing a bundle of positive news can off set the negative impact of announcing a data breach, and Wang et al. [111] show reactions are less damaging when victim firms commit to "action-oriented" security improvements following a breach.

Markets do also reward pre-incident interventions like displaying cybersecurity awareness [8] or obtaining certifications [27, 78], although the impact of certifications is contested [71]. If the certification result holds, it represents an actionable intervention for our hypothetical decision maker. However, security certificates contain noise as well as signal. For example, Rahaman et al. [83] show 86% of the websites in their sample violate the credit card security standard they are certified to. Certified firms perform much better in a similar study of Android apps [69], in which over 98% were compliant.

Legal studies. Studies of legal outcomes for cyber incidents are relatively rare. Romanosky et al. [86] study the likelihood of being sued following a data breach and the proportion of lawsuits that are settled outside of court. They discover that firms offering free credit monitoring after a breach are 6 times less likely to be sued. This represents the clearest evidence regarding the efficacy of a cybersecurity intervention, albeit a post-breach one. The study also provides evidence about prioritising protection efforts given that law suits are 6 times more likely when the breached data contains financial information.

Kesan and Zhang [56] study similar research questions to the earlier study but using data extracted from a proprietary repository. Unlike the credit monitoring finding [86], the explanatory variables in the model[1] do not speak to the efficacy of any actionable intervention. Curiously, the finding that "small companies overall have a higher litigation probability than large ones" [56] flips the sign of the relationship between company size and data breach frequency [10, 113]. A speculative explanation is that larger companies have more mature legal advice regarding interventions to reduce litigation risk. One of the largest cyber insurers supports this view: "there is only an 18 percent chance that a third-party liability action will be brought against one of its customers if

[1] "incident type, whether or not there is the loss of personal sensitive information, number of breached records, company size, and whether or not the company is publicly traded"

one of Chubb's vetted breach response partners is involved, compared to an industry standard of 42 percent" [117]. It should be noted this is far from a sound statistical design given policyholders are not randomly assigned to insurers.

Beyond the US, a study of GDPR fines [88] relied on a third-party to aggregate fines. The authors discovered that Article 32, which concerns data security measures, was amongst one of the three most frequent referenced articles of the GDPR. Ceross and Simpson [21] use freedom of information requests to study fines issued by the United Kingdom's data commissioner. Neither study finds evidence regarding the efficacy of actionable interventions.

Summary. This section surveyed evidence about interventions reducing cyber risk, drawing on surveys from computer security [120] and finance [73]. In general, there is little quantitative evidence that can guide cybersecurity decisions. Technical studies either focus on relatively narrow outcomes like the compromise of web servers or use statistical models that cannot isolate the effect of individual controls [120]. Finance studies have not linked ex-ante security measures to firm-level outcomes like the probability or impact of a cyber incident [73], although stock markets seem to reward security investments like obtaining certifications.

There is stronger evidence regarding ex-post interventions. Corporate officers control information flows in order to mitigate the severity of stock market reactions [3, 39, 111]. The most effective intervention seems to be offering free credit monitoring following a breach [86]. Admittedly, this section has focused on a relatively narrow form of evidence, specifically whether a binary variable describing the intervention provides explanatory power over cyber risk outcomes. The next section describes alternative methods of generating knowledge.

3 How Quants, Techies and Lawyers Generate Knowledge

Although reduced form statistical tests could guide decisions without any understanding of the system, the three disciplines can instead draw on knowledge about the underlying phenomena. For technicians, this relates to the design and implementation of computer systems. Financial professionals might instead reason about the mechanisms and beliefs driving markets. Lawyers can study the internal logic of the legal system. We sketch how each reasons in turn.

3.1 Computer Security Reasoning

The production of knowledge about the security of computer systems is often discussed as the "science of security". In surveying the various positions within the debate, Herley and van Oorschot [48] argue that "practices on which the rest of science has reached consensus appear little used or recognized in security". For example, the authors suggest "observations demonstrating improved outcomes" [48, p. 12] are rare. Verendel confirms that there is little evidence about the effectiveness of security measures in his 2009 meta-review [108]. The argument is supported by the papers [48, 108] described in the previous section.

Concluding that the lack of such findings represents a failure of security as science reflects a narrow philosophy of science [97, p. 2]. Security research has alternative ways of generating knowledge. However, we argue that the knowledge base is continually eroded by technological flux and the bar it must reach rises constantly due to increasing complexity. We will later argue legal knowledge generation avoids these issues.

Formal verification represents a mainstream form of knowledge creation, which involves proving properties of mathematical models that supposedly represent real-world systems. This involves understanding a model via inductive reasoning. Knowledge about security increases with each additional proof. Conducting such analyses at design time can even guide design choices, as with version 1.3 of the TLS protocol [24]. Advances in the tools and concepts used for formal verification also improves the meta reasoning power of the field. This undoubtedly leads to ever increasing certainty about the set of mathematical models used to reason about real-world systems.

This certainty does not always translate into real-world guarantees because the mapping from model to real system is imperfect. For example, researchers did not model how increasing password complexity would change user behaviour and introduce security issues that the model could not express [48]. This problem is even worse for evaluating firm-level security outcomes because "we cannot verify the security of all interactions" [52] between sub-components even if each sub-component is formally verified in isolation. Barring a major break-through, formal verification will not guide cybersecurity decisions at the firm level any time soon.

A more practical approach is to use past attacks to guide decisions. This is how knowledge accrues to the InfoSec "oracles and soothsayers" [41]. A common approach is patch management in which vulnerabilities in software are discovered or observed in the wild, fixes are developed, and then applied by firms who deploy the software. Again, knowledge about possible attacks increases with each new vulnerability. Empirically this can be seen in the proliferation of CVE IDs over time [98].

These incremental knowledge increases may be simply swamped by software complexity [40, 84], as was the case with formal verification. Schneier [89] poses the question in terms of whether vulnerabilities are sparse or dense. In a sparse world, the application of a given patch does not change much as the attacker can exploit other vulnerabilities. Even in a dense world where all vulnerabilities could be discovered with enough time, firms moving to new software versions and systems restarts the process. It has been shown that although reports dry up for individual bug bounty programs, new programs mean hackers still report new bugs [99]. Another aspect undermining knowledge is that attackers can shift to entirely new aspects of the system, such as when attackers discover new classes of attack [58].

Granted, this section considers only two approaches to generating knowledge. Rather than spell out the same argument for each alternative approach, we suggest that a combination of the same considerations—complexity, technological flux, active adversaries and so on—will over-power any knowledge generation that takes place by reasoning about the design of technical systems. We now turn to how quants generate knowledge.

3.2 Quant Reasoning

Again, there are many ways to generate knowledge via quantitative reasoning. We focus on the outlook for the first two approaches identified in the previous section.

Technical Studies. Our survey failed to identify studies isolating the causal effect of individual controls, but will this change? Randomised Control Trials (RCT) are seen as the highest standard of evidence in public health [19]. Indeed, security RCTs have delivered promising results—notifications about open vulnerabilities and compromised assets have been shown to cause improvements in time to patch [100] and clean-up times [63, 107] respectively.

Again, we face the problem of narrow/small pieces of knowledge about complex systems. The RCTs show that notifying system owners about CVE-X causes CVE-X to be fixed sooner and more often, but this may simply displace the attacker's attention to CVE-Y. Thus, we do not know whether such notifications produce better security outcomes at the system level. Public health authors have also argued that RCTs may not be appropriate for generating knowledge about complex systems [109].

This highlights a trade-off between the number of observations and the level at which the system is being observed. For example, Soska and Christin [96] build a data-set of five million websites to study which websites are compromised, meanwhile studies of the financial cost of cyber incidents number in the hundreds [85]. A hypothetical decision maker is likely much more interested in the financial cost of a cyber incident as compared to whether a web page will be compromised. Complexity presents a different problem here—modern firms are so complex that firm-level losses are relatively infrequent and this undermines statistical power.

Machine learning studies [12, 66] represent another way in which complexity limits knowledge generation, here the complexity of the mathematical representation. Such studies aim to optimise prediction rates, often by throwing more explanatory variables into the model. While feature importance analysis can isolate the most important explanatory variables, the statistical relationship is hard to interpret. For example, Liu et al. find that "untrusted HTTPS is by far the most important" [66, p. 1019] feature but its contribution to the prediction is dependent on the firm's other features.

In considering two of the most exciting research designs, namely RCTs and machine learning representations, we are not optimistic about quantitative evidence guiding cybersecurity decisions in the future. While knowledge can be reliably generated about sub-components, especially web infrastructure that is easily scanned [66, 101, 106], system level observations are relatively scarce and also commercially sensitive. This is not true of financial data like stock market price, which is continuously updated.

Finance Studies. Geer et al. [41] suggested portfolio theory and insurance could provide a blueprint for quantitative cybersecurity. Portfolio theory concerns the movement of asset prices over time, which do seem to respond to announcements of security measures [8, 27]. But are markets generating knowledge? Such studies draw on the efficient markets hypothesis [34], which holds that asset prices reflect all publicly available information. The theory implies that the change in stock price following a security investment (or a breach) reflects investors expectations about the increase (decrease) in the firm's future earning potential. Putting aside the role of information distortions [3, 39, 64], such changes are based on the investors' expectation that firms with a given security investment are exposed to less risk. This expectation cannot be derived a priori and therefore must have come from some alternative information source. Providing the expectation is valid, the alternative information source generated the knowledge, not the market.

More generally, portfolio theory abstracts away from the kind of security decisions we are interested in. Power [82] suggests quantitative models are in conflict with pragmatists interested in "internal controls" [74]. This is confirmed by a survey [120, Tab. III] of studies of stock market prices that shows such studies tend not to consider the role of preventative security.

Insurance is different because insurers collect information about internal controls during the application process [76, 87, 121], which could be linked to the financial cost of cyber incidents via insurance claims. This system-level outcome variable means that discovering internal controls with explanatory power would represent knowledge that could guide decisions (although identifying causality may be difficult without a source of randomness). So far, insurers have experienced problems like inconsistent data collection and an unwillingness to share data [121]. Admittedly, it would be premature to rule out the possibility of insurers generating knowledge but we see few signs for optimism.

3.3 Legal Reasoning

The previous section showed research has not *yet* generated quantitative evidence about cybersecurity decisions. The previous two subsections argued that doing so is unlikely in a series of arguments from first principles. Broadly, we argued that systems complexity, technological flux, and strategic attackers would erode the knowledge base before it was sufficiently advanced to guide decisions.

This subsection argues this is different for legal reasoning. We argue that there are multiple ways in which the law resolves uncertainty through its openness to interpretation and application. This characteristic operates: (a) within the nature of law and architecture of its application (e.g. the courts); (b) the language of the law; (c) establishing legal certainty through (i) courts and (ii) the lawyer providing advice to a client. We address each in turn.

The nature of law and legal systems. Fundamentally, the law consists of formalised rules intended to regulate behaviour. While there are different legal traditions and architectures (e.g. the common law and civil law traditions), they all share common elements.[2] These rules may be divided into two broad types: (i) legislation, that is, the explicit, stated rules promulgated by a sovereign legislator and (ii) the decisions of the judiciary. The former provides the hard rules to be applied and the latter informs the norms that guide application to a particular set of facts. Legal systems are inquisitive and embody an adversarial process in order to establish accepted facts and applicable law. In effect, there are generally three entities participating in any given legal action: (i) the complainant, the party bringing a claim (ii) the defendant, the party against whom the claim is brought and (iii) the judiciary, who officiates and regulates the process. However, the law is much more than mere rules as it embodies the culture and society in which it operates [105]. This is evident in the pronounced manner by which law relies on its past and tradition as a source of authority for determining problems presented in the present [60].

The corpus of tradition is applied as a matter of *interpretation*, which gives acknowledgement of the past and brings it into the present for evaluation and discernment. In this way, applications of the law may be viewed as exercises in hermeneutics, attempting to bring the uncertainty experienced in the present in continuity with the authoritative past. Of course, this does not preclude that interpretation may introduce novelties and disapply previously held positions. Nonetheless, the effect of the emphasis on tradition is that it strengthens the ability of law to enact social control. Such control does not lay in the abstraction of principle nor strict dogmatic adherence to previous articulations of said principles, but rather in its pragmatic approach to provide effective solutions to real world problems [81]; the law can be said to be based in such unashamed practicality, that Posner argues much legal theory scholarship is underdeveloped to such an extent that it is ultimately "vacuous" [80, p.3].

Interpretation does not merely apply to pronouncements of the courts. Lawyers advise their clients with a strategy for navigating legal obligations and similar requirements. The advice provided by the lawyer in such scenarios is not only informed by the individual understanding of the law (which includes knowledge of the rules themselves and how they have been applied in the past), but also the individual's personal and professional morals and ethics [110]. Legal advice and strategies are therefore variable to a degree and such advice may have not been tested in a court, which would evaluate its merits. This is often the case with novel, rapidly developing areas of law, especially those involving technology [13].

Law and language. Language is the medium by which law operates, which has its benefits and drawbacks in knowledge generation and the management of legal actions. The language used in the law (both legislation and judicial proceedings) is characterised by its dense and often archaic sentence structure, unusual definitions for common terms, and foreign vocabulary (e.g. Latin and Norman French in the case of the English legal tradition) [70]. The language of law is to be navigated via rules of interpretation which may not be readily apparent to the intended regulated entities, which necessitates 'translation' by trained legal professionals [43]. HLA Hart [47] viewed this as a feature, highlighting that it is through the law's "open texture" that we can apply it effectively to the diversity of human social situations requiring adjudication and remedy. In more practical terms, legislation aimed at regulating cybersecurity may not be immediately comprehensible to those operating in this field. For example, the rules on 'privacy-by-design' for data protection have been criticised by both lawyers and engineers for its imprecise language and difficulty in application to real-world situations [46, 115]. On the other hand, the law may be more able to provide a solution to a dispute despite the high levels of complexity involved in characterising a case.

The law is also not an autopoietic system because it relies on the input of interpretation in order to self-perpetuate [68]. Distrust of over-extending definitions and argumentation has been found to be common in law [67], which may act to reinforce consistency of terminology and reasoning. However, there are instances wherein the legal process may reframe and re-define previously established concepts to the extent that it represents a new paradigm for a field within the law. Furthermore, there are instances where the intended scope and meaning of the legislation is not semantically or conceptually congruent with the actuality of the situation it addresses [114]. This may be due to lack of understanding on the part of the legislator or because consistency with

[2] For the purpose of this paper, we treat the elements of legal architecture in a singular manner, without discussion of the distinctions and differences in various traditions and jurisdictions. While these distinctions are meaningful, exploring this nuance detracts from the overall argument being made about the nature of law in contrast to technical approaches to cybersecurity. That said, the authors are undoubtedly more influenced by the common law tradition.

other elements of legislation required such linguistic construction. For example, technology law is often intended to be 'technologically neutral' in order to future proof the legislation against new developments. Regardless, there are circumstance that stretch the limits of legal interpretation. A teleological approach may be applied to bridge such a divide, although judges are cautious to not accept overly ambitious interpretations which can distort and denature the meaning of the words to the point of uselessness [2].

Establishing legal certainty. From the above, the law may be seen as fluid and flexible. We argue that fluidity is the law's strength when adhering to its function of practicability and consistency in generating binding decisions. While this allows the law to adapt to real-world circumstances, it also opens a vulnerability to inconsistent interpretation. When considering how to establish certainty in the legal process, Braithwaite [16] distinguishes between the use of 'rules' and 'principles'. In circumstances which are considered in law to be stable and without high stakes, the judicial process employs well-defined 'rules'. However, as discussed, a lawyer may successfully argue for a novel interpretation of the rules which generates uncertainty for the legal understanding of a particular phenomenon. Therefore, Braithwaite suggests, the law retreats to a set of 'principles' which provide boundaries to the rules in order to avoid interpretations that defeat the original purpose of the law. In this way, the legal process is able to circumvent the worst effects of creative interpretation, although this is not always the case.

In the law, attempts at reasoning about a problem according to the legal tradition often give way to analogy. The use of analogy may act as a short-hand for difficult and complicated reasoning. However, MacCormick [68] argues that analogy is to underscore the principles to be applied. Hunter [51] provides a comprehensive overview of the legal theory, bolstered by cognitive science, that analogy use in precedent has a two-process structure: discovery and justification. In the first step, the adjudicator identifies the appropriate outcome to the problem and then provides the reason ('justification') for reaching this conclusion in the second.

Thus, while there is much to critique in the application of the law, its ability to provide a resolution is to be appreciated. This resolution need only be *acceptable* (or in the worst cases *tolerable*) to the parties to a case and wider society. An acceptable solution does not necessarily mean it is the most appropriate or most considered. It does however mean that when faced with similar sets of facts for which there is a court decision, there is therefore a known outcome around which a cybersecurity strategy may be devised.

This strategy need not be a box-checking exercise related to obligations identified in past decisions, rather it can be a strategic exercise in which the firm actively shapes the law. This is achieved when lawyers adopt a legal strategy that embodies a creative conflict management approach to novel challenges rather than merely managing transactions between the client and other parties [123]. For example, a bank invested in the development of a pro-active, 'best-in-class' corporate governance structure because the regulator was concerned about short-comings in this area. The bank later found that this facilitated approval of a bank's acquisition by a now sympathetic regulator [13, 14]. Instead of viewing governance as a compliance obligation, the bank transformed it into a competitive, commercial advantage.

4 Legalised Cybersecurity

This section sketches the implications of *lawyers inheriting cybersecurity.* Some of these changes have already been realised, while others are mere predictions. We focus on the meta-reasoning through which decisions are made rather than predicting specific outcomes of decisions. In particular, we focus on how legal reasoning moves from a descriptive assessment, such as the statistical tests in quantitative cybersecurity, to one of normative values. The resulting changes span epistemology, work culture, and economics.

Information Control. Technical and legal risk can be compared in terms of the role of information. Technical risk considers how likely an adversary is to use information about technical vulnerabilities to compromise a system. Legal risk considers how likely a complainant or the judiciary is to use legally relevant information against the defendant. Whereas technical risk can be mitigated by sharing information and fixing vulnerabilities, legal risk is amplified when information is shared as it could be used as evidence against the defendant. In this way, the legal view breaks from the scientific perspective that evidence should be shared widely.

We can expect lawyers inheriting cybersecurity to reduce information sharing. Indeed, organisations are less likely to disclose information when they could be held liable for doing so [61]. A less widely appreciated implication is the potential for lawyers to influence the creation and documentation of evidence. Such influence anticipates pre-litigation discovery processes in which claimants and defendants can request evidence from each other (and these can be legally enforced via motions to compel). For example, shareholders may sue a breached company and request the forensic report detailing which security procedures were in place. Lawyers can mitigate this legal risk in multiple ways, which are illustrated when lawyers run incident response (notably thousands of incidents are already run by lawyers each year [118]).

Cyber law firms argue they must be placed at the centre of incident response (as in Figure 1) in order to maintain client-attorney privilege [118]. The breached firm's attorneys can claim evidence is privileged, and therefore not discoverable, providing it was produced *in anticipation of litigation.* To argue this more clearly, law firms hire the forensic providers only after an incident is known. This avoids the situation following Capital One's data breach. A judge ruled that a

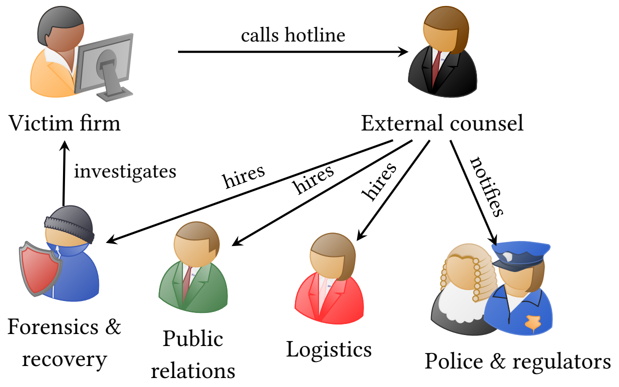

Figure 1. External counsel control the incident response value chain and interactions with authorities.

forensic report was not protected by client-attorney privilege because the contract with the forensics firm was signed before the incident[3].

Client-attorney privilege is never absolute, which motivates a second mitigation strategy of changing information creation. In their study of incident response, Woods and Böhme [119] find that lawyers use "informal/verbal reports to avoid documents that could be discovered by a litigant" and that "investigations are structured according to the desires of the law firm/lawyer". While further research is required to understand the full implications, it is reasonable to conclude that less documentation limits knowledge generation by quants. In this way, lawyers inheriting cybersecurity may actively undermine the science of security.

Reasonable over effective. Risk decisions may be guided by concepts like *reasonableness* or *appropriateness* rather than *effectiveness*. This is the legal equivalent of the wisdom behind "no-one gets fired for buying IBM" [4]. Whereas an effectiveness criterion strives to be better than current practice and rewards innovations that do so, reasonableness emphasises following established practices and may even punish deviating from them. For example, Selzer analyses appropriate data protection controls under the GDPR and finds that it "will not always be easy" [90, p. 123] to identify the "state of scientific and technical knowledge" as the law recommends. Selzer instead suggests that firms should follow "established recommendations" offered by regulators [90, p. 123]. Deviating from such recommendations brings legal risk unless the deviation is based on scientific knowledge, which is broadly unavailable (see Section 2).

Such recommendations look like *security compliance*, which Julisch [55] defines to be "conformance with a given set of security requirements". Julisch argued this the drive for compliance better explained how organisations make decisions than the prevailing concept of *security*, defined as "the state of being safe from threats". He suggests

> "security incidents are tolerated more easily if one can show that they occurred despite the affected IT system being compliant with all applicable security regulations" [55, p. 1]

For example, Thaw finds that "a primary effect of breach notification laws was to focus intensive effort on encryption of portable devices and media containing personal information" [p. 160][104]. This suggests that rather than invest resources in preventing breaches (*security*), organisations invested in reducing the legal risk associated with a breach. Such herding behaviour may increase homogeneity and with it systemic risk [42].

Legacy Concepts. In computer science, legacy systems involve software or hardware for which newer versions are available. While legacy problems are sometimes accepted because of the economic cost of switching, there is a secular force shifting technology adoption towards new code bases, hardware and the next generation of technicians. In contrast, switching to an entirely new contractual language or interpretative framework is profoundly undesirable for lawyers as this would throw away the certainty developed over time.

This can be seen in how lawyers drafted cyber insurance exclusions. Rather than tailor terms and concepts based on which kinds of cyber attack were uninsurable, cyber insurance policies contain war clauses written for the twentieth century [122]. US courts will now have to decide whether the NotPetya cyber attack triggers war clauses based on case law that has been inconsistently tested by conflicts as diverse as "the American Civil War, Pearl Harbour and the 9/11 attacks" [122].

It is, admittedly, not unreasonable that case law might reach far back into history, but even written law can reference legacy technologies. For example, the 2020 Brexit withdrawal agreement mentioned Netscape Communicator as a modern service[4]. While the cybersecurity implications of each example is unclear, the general point remains—lawyers intentionally invoke legacy dependencies, whereas techies try to avoid them. It is worth asking whether legacy concepts and the pursuit of reasonable security measures make defenders less responsive, in contrast to attackers who face no similar constraints.

Sociology of Risk. Beyond the relatively abstract shifts in epistemology and information sharing identified so far, lawyers inheriting cybersecurity will strengthen the norms and traditions of lawyers. Returning to Figure 1, we see that lawyers hire forensics firms or at least recommend which firms to hire. At the margin, this will reward technical professionals who are seen favourably in the eyes of lawyers. Woods and Böhme find that:

[3] https://www.huntonprivacyblog.com/2020/07/07/forensic-report-deemed-not-privileged-capital-one-ordered-to-release-report/

[4] https://www.bbc.com/news/technology-55475433

> "lawyers emphasised the importance of non-technical factors like responsiveness, communications with clients, and a willingness to accept work (e.g. not to refuse incidents and to provide all required services)" [119, p. 13]

Such factors may be at odds with the cultural values of the technical community [94, 102]. It would be premature to draw definitive conclusions, but it suggests areas for future work.

In particular, future work should not frame risk management decision making as a determinate process converging on objective truth [112]. Rather risk decisions are made by self-interested actors who construct "risk objects" in order to gain influence within and over other organisations [49]. Mandatory data breach notification laws represent a very practical example. Although notifying all affected individuals presents itself as a determinate process, different disciplines favour different methods. Lawyers recommend manual analysis via a team of paralegals, which can cost up to $500k [119, p. 21], whereas technical practitioners favour a cheaper approach using probabilistic models to identify personal data in unstructured data breaches (e.g. an email inbox) with the accompanying true negatives and false positives meaning not all individuals are notified [118]. Each process will lead to a different result and adoption will depend on the relative political capital of each practitioner. Qualitative methods [20, 36, 45] can uncover similar practices within cyber risk mangement.

Economics and Business Models. We should also anticipate shifts in both internal business structure and also how businesses transact with each other. For example, the lawyer led model of IR disrupts the integration of ex-ante monitoring and ex-post investigation recommended in NIST-800-61 [23]. In traditional IR, visibility from network monitoring helped to guide investigations, logs were set up to collect evidence, and internal investigators were often familiar with the systems. In the new model [118], lawyers choose the firms they want to work alongside independent of whether the firm has existing network access.

This introduces new business logic. In the old model [23], firms first sold monitoring and offered investigation as an optional follow-on service. In the new model, firms first sell investigations and then try to sell networks tools used during investigation as an on-going service [119]. Interestingly, the "lawyers we spoke to felt this was unprofessional", which contrasts with the security professionals belief that their "duty is to improve the client's security posture wherever possible" [119]. Again, we see potential frictions resulting from the values of each discipline.

In terms of internal business structure, much has been written about how the so-called "tournament of lawyers" [37] created the modern law firm. Here, junior employees compete to rise in performance rankings as the top ranked juniors are rewarded with partnership. A similarly definitive study on the structure of security vendors is not available, but it is worth contrasting with the typical start-up model in which employees are rewarded with more equity the earlier they joined. Again, we make no definitive conclusions and remain open to the possibility that cyber law firms will adopt the start-up model. For example, Mullen Coughlin, founded in 2016, "disrupted" more traditional law firms in becoming the dominant cyber insurance law firm [119, Fig. 7].

Summary. We do not make precise predictions about the impact of lawyers inheriting cybersecurity, rather we suggest a number of structural changes in how decisions are made. The concept of *effective* security controls is rooted in the scientific method's ability to quantify effect sizes, but it will soon be replaced by notions like *reasonable* or *appropriate* controls. This results from the failure of quants to produce evidence about effectiveness (see Section 2). Evidence generation will be further undermined as lawyers' risk-averse information control displaces the commitment to open data found in both the scientific and security community. We also suggest that technologists' drive to innovation will be tempered by the law's inclination to import the conceptual certainty found in precedents and case law.

Whereas changes in meta-reasoning behind decisions are abstract and difficult to test empirically, social science methods are well-placed to test the existence and impact of shifts in the cultural, economic and organisational context in which cybersecurity decisions are made. Researchers could begin to ask questions like: If law schools "turn students into lawyers" [38], have lawyers identified an equivalent institution producing security professionals? Does this come at the cost of uncredentialed professionals? Do law firms recognise the synergies between ex-ante monitoring and ex-post investigation, or are they instead focused on the resulting conflicts of interest? Will security vendors adopt "tournament theory" [37] as lawyers have done, or will lawyers internalise tech's disruptive spirit? Such questions are natural to ask within the "lawyers shall inherit cybersecurity" paradigm.

5 Discussion

This raises the question of what we mean by *inherit.* It is broader than more cybersecurity law and policy coming into effect. Rather, our argument centres on the hierarchy of power and meta-reasoning through which cybersecurity decisions are made in firms. *Lawyers inheriting cybersecurity* means that legal reasoning—referencing cybersecurity law, regulatory guidance, precedents and the anticipatory interpretations of individual lawyers—will displace reasoning about security as a technical property. This section discusses progress towards the inheritance, who is driving it, and how the technical community will be integrated into the new decision-making hierarchy.

Progress to Date. Comprehensively answering the question *to what extent does legal reasoning already drive decisions* would require a dedicated study. However, a number of principles we have outlined can already be seen, such as the shift from effective to reasonable security measures. In the United States, there is an emerging body of law on what constitutes "reasonable" security in the eyes of the states [93], the FTC [17] and the SEC [59, Chapter 4]. It is an open question as to whether firms look to such guidance when making cybersecurity decisions. As mentioned earlier, information control and the economic effects can be seen in cyber incident response [118]. One technical firm reporting that 50% of their engagements are conducted under the direction of an attorney [25]. We can also look to contractual agreements between firms as a form of legalisation. The PCI DSS requirements [75] influence data security practices when it comes to handling credit card data [69, 83].

Despite judgements being made and new laws coming into effect, regulatory attention and capacity seems to be lagging cybersecurity incidents. Even though the absolute number of private and public data breach actions brought in federal courts steadily increased "from a couple of dozen in 2005, to almost 200 by 2014" [85], this is still an order of magnitude lower than the number of cyber incidents in that time [120]. GDPR enforcement rates have been similarly underwhelming [88, 116]. Still, enforcement is not needed when 68% of organisations believed that the GDPR would "dramatically increase the costs of doing business" [103] and likely sought out legal advice, either internally or on the market.

Enforcement rates also vary by organisational characteristics; a greater proportion of large firms were fined under the GDPR than small or medium firms [91], although Kesan and Zhang find the opposite relationship in the US [56]. The influence of lawyers will track the attention of regulators, both in terms of past enforcement and anticipated actions. Together, these findings suggest the inheritance is in progress, which raises the question of who is driving it.

A Reluctant Inheritance. While inheritance may conjure an image of greedy children fighting over the family jewels, the reality can be more like children manoeuvring to avoid the responsibility of running the family business. No doubt there are ambitious lawyers and law firms fighting to win as much work as possible, much like the greedy children and the family jewels. But is the wider discipline actively sidelining quants and techies? Here, the children trying to avoid responsibility seems more apt.

It is instructive to return to Article 32 of the General Data Protection Regulation. The law willingly defers authority to the quants by recommending that technical and organisational measures be identified and selected based on the "state of scientific and technical knowledge" [90]. However, the dearth of evidence forced the law's pragmatic drive for certainty to reluctantly to take over. This manifested as the national regulators issuing recommendations, to which organisations are now recommended to turn [90]. Thus, we argue law is reluctantly inheriting cybersecurity.

The reluctance means the law will not eliminate all ambiguity, in fact it falls far short of this. Residual uncertainty results from cases where the law is not able to rely on established precedent in order to reason, such as where the problem encountered is novel or there have not been many cases [2]. There is further uncertainty in legal issues where the law may perhaps over prescribe actions without filling in the details. This is arguably the case in data protection law wherein technical measures are alluded to without the explicit specification necessary for implementation [22, 91].

Influencing Legalised Cybersecurity. This reluctance creates space for the technical and scientific community to still influence cybersecurity decisions within the legalised cybersecurity paradigm. The "oracles and soothsayers" [41] cannot rely on firms recognising their received wisdom. Instead security practitioners must contribute to cybersecurity standards development or begin to document their knowledge, a long standing problem. For example, the Ohio Data Protection Act creates a liability safe harbour for firms who implement recognised cybersecurity standards and frameworks [93, p. 22]. Similarly, the GDPR tells firms to look to "state of scientific and technical knowledge" [90]. In this way, the inheritance disrupts the established hierarchies and establishes new hierarchies, with power annointed according to how easily expertise is recognised by the legal system.

An alternative form of influence is rejecting the authority of lawyers. A quiet revolution that will surely take place is "shadow security" [57] in which employees interpret firm-wide security policies with a view to maintaining business function. For example, an organisation responding to Not-Petya abandoned communications protocols intended to maintain client-attorney privilege because it was deemed more important to restore business function than mitigate litigation risk.[5]

Finally, the tide could turn against lawyers if the growth of technical risk out-paces that of litigation risk. For example, lawyer-led incident response was clearly beneficial when the primary cost-driver was post-breach litigation. The added-value is less clear when "litigation rates are around 1% while ransomware payments grow 1000% year-on-year" [118]. Such a movement will be tempered by the tendency for power hierarchies to entrench themselves over time.

Similar Theories. We are neither the first nor likely the last to argue that law will influence a problem area with increasing societal importance, nor to make this argument specific to cybersecurity. Given the United States' influence

[5] Based on correspondence with a senior employee involved in the response.

over cybersecurity, the general point goes back to De Tocqueville's 19th century writings [26] and the common sense wisdom behind the phrase 'a nation of lawyers'. The specific argument can be found in many forms.

Lessig [62] identifies that given extra-legal forces influence online behaviour, public policy should shape each of markets, norms, and computer code. We concur with Lessig by recognising that legal compliance cannot unilaterally determine cybersecurity decisions, rather legal reasoning will provide the decision making structure through which decisions are made. In this sense, we document the realisation of Lessig's normative argument, while also questioning the extent to which law has and will positively impact cybersecurity outcomes.

We also provided a wider perspective than Julisch [55] who emphasised the compliance aspect of legalisation, or Anderson [5] who emphasised liability dumping. While both aspects are important under certain circumstances, the compliance perspective cannot explain the ambiguous areas like privacy engineering in which laws reference technological goals without enough specificity to guide technical decisions. Further, both perspectives miss the cultural and business aspects of lawyers influencing cybersecurity decisions.

6 Conclusion

For the last 20 years, various authors have argued that the status quo "of oracles and soothsayers" [41] making InfoSec decisions based on experience and intuition will be replaced by a new approach guided by scientific evidence. The "InfoSec belongs to the quants" paradigm recommends that firms build security programs based on the effectiveness of individual controls.

The first part of our paper showed that the evidence base to put quantitative security into practice is lacking. This claim is based on evaluating available empirical evidence identified in various surveys [73, 108, 120]. Broadly, we showed that complexity, technological flux, and strategic adversaries undermine knowledge generation via quantitative cybersecurity. For example, the gold standard of scientific evidence, randomised control trials, are yet to provide results about system-level outcomes. We made similar arguments about why methods like stock market reactions, formal methods, and machine learning representations will not generate knowledge to guide firm-level decisions any time soon.

Our second contribution argues that cybersecurity decisions will increasingly be guided by legal reasoning. Whereas uncertainty is compatible with quantitative cybersecurity in that statistical tests can indefinitely produce null results, law as a system of social control can only tolerate so much ambiguity. Various legal mechanisms function to guide firms in mitigating the legal aspects of cyber risk. This ranges from courts making judgements about the laws, regulators issuing guidance, and individual lawyers advising firms thereby establishing common practice. The *lawyers will inherit cybersecurity* paradigm centres on how the ability to reason about legal mechanisms provides relatively more certainty than the equivalent processes for reasoning about technical risk.

We can only speculate about the impact of this development. For example:

- Controls will be prioritised based on *reasonableness* or *appropriateness* rather than *effectiveness*.
- Information will be less widely shared.
- Security will absorb the professional culture of lawyers by osmosis.
- Business models and firm structures that match or look familiar to law firms will be rewarded and so on.

We then cautioned against misreadings of our paradigm. The law is less seizing control of cybersecurity and more reluctantly taking authority due to the knowledge vacuum. As a result, law leaves many problems ambiguous. One can expect different areas of cybersecurity to be more or less exposed to law as a system of social control, and our paradigm's explanatory power tracks these variations.

As such, we recommend a research program understanding the law as both an abstract system of rules and also a social system of actors [32]. The rules can be understood via legal reasoning, whereas qualitative methods are needed to understand how practitioners interpret and operationalise the rules. Understanding the impact of *lawyers inheriting cybersecurity* requires a range of lenses including legal, technical, economic, sociological and more.

Acknowledgments

We thank the NSPW participants for a lively and fruitful discussion, in particular Rikke Bjerg Jensen and Simon Parkin for acting as pre and post workshop shepherds. We also thank all the Incident Response practitioners who extended their time to help an ignorant researcher understand the ecosystem—hearing your perspecives brought the first author to this idea. The first author is funded by the European Commission's call H2020-MSCA-IF-2019 under grant number 894700.

References

[1] Inaki Aldasoro, Leonardo Gambacorta, Paolo Giudici, and Thomas Leach. 2020. The drivers of cyber risk. *BIS Working Paper* (2020).

[2] Larry Alexander. 1997. Banality of legal reasoning. *Notre Dame Law Review* 73 (1997), 517–534.

[3] Eli Amir, Shai Levi, and Tsafrir Livne. 2018. Do firms underreport information on cyber-attacks? Evidence from capital markets. *Review of Accounting Studies* 23, 3 (2018), 1177–1206.

[4] Ross Anderson. 1993. Why cryptosystems fail. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*. 215–227.

[5] Ross Anderson. 1994. Liability and computer security: Nine principles. In *European Symposium on Research in Computer Security*. Springer, 231–245.

[6] Ross Anderson and Tyler Moore. 2006. The economics of information security. *Science* 314, 5799 (2006), 610–613.

[7] Wade H Baker, Loren Paul Rees, and Peter S Tippett. 2007. Necessary measures: metric-driven information security risk assessment and decision making. *Commun. ACM* 50, 10 (2007), 101–106.

[8] Henk Berkman, Jonathan Jona, Gladys Lee, and Naomi Soderstrom. 2018. Cybersecurity awareness and market valuations. *Journal of Accounting and Public Policy* 37, 6 (2018), 508–526.

[9] Peter L Bernstein. 1996. *Against the gods: The remarkable story of risk.* Wiley New York.

[10] Claudia Biancotti. 2018. The price of cyber (in)security: evidence from the Italian private sector. In *Workshop on the Economics of Information Security.*

[11] Christian Biener, Martin Eling, and Jan Hendrik Wirfs. 2015. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance-Issues and Practice* 40, 1 (2015), 131–158.

[12] Leyla Bilge, Yufei Han, and Matteo Dell'Amico. 2017. Riskteller: Predicting the risk of cyber incidents. In *Proc. of the Conference on Computer and Communications Security.* ACM, 1299–1311.

[13] Robert C Bird. [n.d.]. Pathways of legal strategy. *Stanford Journal of Law, Business & Finance* ([n. d.]).

[14] Robert C Bird. 2011. Law, strategy, and competitive advantage. *Connecticut Law Review* 44 (2011), 61–97.

[15] Antoine Bouveret. 2018. Cyber risk for the financial sector: a framework for quantitative assessment. In *IMF Working Papers (43/145).*

[16] John Braithwaite. 2002. Rules and principles: A theory of legal certainty. *Australian Journal of Legal Philosophy* 27 (2002), 47.

[17] Travis D Breaux and David L Baumer. 2011. Legally "reasonable" security requirements: A 10-year FTC retrospective. *Computers & Security* 30, 4 (2011), 178–193.

[18] Ryan Calo. 2019. Privacy Law's Indeterminacy. *Theoretical Inquiries in Law* 20 (2019), 33–52.

[19] Nancy Cartwright and Jeremy Hardie. 2012. *Evidence-based policy: A practical guide to doing it better.* Oxford University Press.

[20] Myriam Dunn Cavelty. 2018. Cybersecurity research meets science and technology studies. *Politics and Governance* 6, 2 (2018), 22–30.

[21] Aaron Ceross and Andrew Simpson. 2017. The use of data protection regulatory actions as a data source for privacy economics. In *International Conference on Computer Safety, Reliability, and Security.* Springer, 350–360.

[22] Aaron Ceross and Andrew Simpson. 2018. Rethinking the Proposition of Privacy Engineering. In *Proceedings of the New Security Paradigms Workshop* (Windsor, United Kingdom) *(NSPW '18)*. 89–102.

[23] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. 2012. Computer security incident handling guide. *NIST Special Publication* 800, 61 (2012), 1–147.

[24] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. 2017. A comprehensive symbolic analysis of TLS 1.3. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security.* 1773–1788.

[25] CrowdStrike. 2020. CrowdStrike Services Cyber Front Lines Report!: Available: https://go.crowdstrike.com/rs/281-OBQ-266/images/Report2020CrowdStrikeServicesCyberFrontLines.pdf. [Online; accessed 24-Dec-2020].

[26] Alexis De Tocqueville. 1835. *Democracy in America.* Vol. 1. Saunders and Otley (London).

[27] Jason K Deane, David M Goldberg, Terry R Rakes, and Loren P Rees. 2019. The effect of information security certification announcements on the market value of the firm. *Information Technology and Management* 20, 3 (2019), 107–121.

[28] Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. 2016. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity* 2, 1 (2016), 3–14.

[29] Benjamin Edwards, Jay Jacobs, and Stephanie Forrest. 2019. Risky business: Assessing security with external measurements. *arXiv preprint arXiv:1904.11052* (2019).

[30] Martin Eling and Nicola Loperfido. 2017. Data breaches: Goodness of fit, pricing, and risk measurement. *Insurance: Mathematics and Economics* 75 (2017), 126–136.

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

[32] Howard Erlanger, Bryant Garth, Jane Larson, and Elizabeth Mertz. 2005. Is it time for a new legal realism. *Wisconsion Law Review* (2005), 335–365.

[33] Gregory Falco, Martin Eling, Danielle Jablanski, Matthias Weber, Virginia Miller, Lawrence A Gordon, Shaun Shuxun Wang, Joan Schmit, Russell Thomas, Mauro Elvedi, et al. 2019. Cyber risk research impeded by disciplinary barriers. *Science* 366, 6469 (2019), 1066–1069.

[34] Eugene F Fama. 1970. Efficient capital markets a review of theory and empirical work. *Journal of Finance* 25, 2 (1970), 383–417.

[35] Dinei Florêncio and Cormac Herley. 2013. Sex, lies and cyber-crime surveys. In *Economics of Information Security and Privacy III.* Springer, 35–53.

[36] Damjan Fujs, Anže Mihelič, and Simon LR Vrhovec. 2019. The power of interpretation: qualitative methods in cybersecurity research. In *Proceedings of the 14th International Conference on Availability, Reliability and Security.* 1–10.

[37] Marc Galanter and Thomas Palay. 1994. *Tournament of lawyers: The transformation of the big law firm.* University of Chicago Press.

[38] Bryant G Garth and Joanne Martin. 1993. Law schools and the construction of competence. *Journal of Legal Education* 43, 4 (1993), 469–509.

[39] Sebastien Gay. 2017. Strategic news bundling and privacy breach disclosures. *Journal of Cybersecurity* 3, 2 (2017), 91–108.

[40] Dan Geer. 2015. For good measure: The undiscovered. *; login:: the magazine of USENIX & SAGE* 40, 2 (2015), 50–52.

[41] Dan Geer, Kevin Soo Hoo, and Andrew Jaquith. 2003. Information security: Why the future belongs to the quants. *IEEE Security & Privacy* 1, 4 (2003), 24–32.

[42] Dan Geer, Eric Jardine, and Eireann Leverett. 2020. On market concentration and cybersecurity risk. *Journal of Cyber Policy* 5, 1 (2020), 9–29.

[43] John Gibbons. 1999. Language and the law. *Annual Review of Applied Linguistics* 19 (1999), 156–173.

[44] Lawrence A Gordon and Martin P Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* 5, 4 (2002), 438–457.

[45] Evert Gummesson. 2000. *Qualitative methods in management research.* Sage.

[46] Seda Gürses, Carmela Troncoso, and Claudia Diaz. 2011. Engineering privacy by design. *Computers, Privacy & Data Protection* 14, 3 (2011), 25.

[47] Herbert Lionel Adolphus Hart. 2012. *The concept of law* (third ed.). Oxford University Press.

[48] Cormac Herley and Paul C Van Oorschot. 2017. Sok: Science, security and the elusive goal of security as a scientific pursuit. In *2017 IEEE Symposium on Security and Privacy (SP).* IEEE, 99–120.

[49] Stephen Hilgartner. 1992. The social construction of risk objects: Or, how to pry open networks of risk. *Organizations, Uncertainties, and Risk* (1992), 39–53.

[50] Douglas W Hubbard and Richard Seiersen. 2016. *How to measure anything in cybersecurity risk.* John Wiley & Sons.

[51] Dan Hunter. 2001. Reason is too large: Analogy and precedent in law. *Emory Law Journal* 50 (2001), 1197–1264.

[52] Nathaniel Husted and Steven Myers. 2014. Emergent Properties & Security: The Complexity of Security as a Science. In *Proceedings of the 2014 New Security Paradigms Workshop.* 1–14.

[53] Jeevan Jaisingh and Jackie Rees. 2001. Value at risk: A methodology for information security risk assessment. In *Proceedings of the INFORMS Conference on Information Systems and Technology*. Citeseer, 3–4.

[54] Erland Jonsson. 1998. An integrated framework for security and dependability. In *Proceedings of the 1998 workshop on New security paradigms*. 22–29.

[55] Klaus Julisch. 2008. Security compliance: the next frontier in security research. In *Proceedings of the 2008 New Security Paradigms Workshop*. 71–74.

[56] Jay P Kesan and Linfeng Zhang. 2021. When Is A Cyber Incident Likely to Be Litigated and How Much Will It Cost? An Empirical Study. In *Invited Contribution to Symposium on Cyber Insurance, Connecticut Insurance Law Journal, Forthcoming*.

[57] Iacovos Kirlappos, Simon Parkin, and M Angela Sasse. 2015. " Shadow security" as a tool for the learning organization. *ACM SIGCAS Computers and Society* 45, 1 (2015), 29–37.

[58] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, et al. 2019. Spectre attacks: Exploiting speculative execution. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1–19.

[59] Jeff Kosseff. 2019. *Cybersecurity law*. John Wiley & Sons.

[60] Martin Krygier. 1986. Law as tradition. *Law and Philosophy* 5, 2 (1986), 237–262.

[61] Stefan Laube and Rainer Böhme. 2017. Strategic aspects of cyber risk information sharing. *ACM Computing Surveys (CSUR)* 50, 5 (2017), 1–36.

[62] Lawrence Lessig. 2009. *Code: And other laws of cyberspace*. ReadHowYouWant.com.

[63] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remedying web hijacking: Notification effectiveness and webmaster comprehension. In *Proceedings of the 25th International Conference on World Wide Web*. 1009–1019.

[64] Zhaoxin Lin, Travis RA Sapp, Jackie Rees Ulmer, and Rahul Parsa. 2019. Insider trading ahead of cyber breach announcements. *Journal of Financial Markets* (2019), 100527.

[65] Bev Littlewood, Sarah Brocklehurst, Norman Fenton, Peter Mellor, Stella Page, David Wright, John Dobson, John McDermid, and Dieter Gollmann. 1993. Towards operational measures of computer security. *Journal of Computer Security* 2, 2-3 (1993), 211–229.

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

[67] Eric Lode. 1999. Slippery slope arguments and legal reasoning. *California Law Review* 87 (1999), 1469–1544.

[68] Neil MacCormick. 1993. Argumentation and interpretation in law. *Ratio Juris* 6, 1 (1993), 16–29.

[69] Samin Yaseer Mahmud, Akhil Acharya, Benjamin Andow, William Enck, and Bradley Reaves. 2020. Cardpliance: PCI DSS Compliance of Android Applications. In *29th USENIX Security Symposium (USENIX Security 20)*. 1517–1533.

[70] Yon Maley. 1987. The language of legislation. *Language in society* (1987), 25–48.

[71] Dennis D Malliouris and Andrew C Simpson. 2019. The stock market impact of information security investments: The case of security standards. In *Workshop on the Economics of Information Security*.

[72] Mohammad Hossein Manshaei, Quanyan Zhu, Tansu Alpcan, Tamer Bacşar, and Jean-Pierre Hubaux. 2013. Game theory meets network security and privacy. *ACM Computing Surveys (CSUR)* 45, 3 (2013), 1–39.

[73] Michael McShane, Martin Eling, and Trung Nguyen. 2021. Cyber risk management: History and future research directions. *Risk Management and Insurance Review* (2021).

[74] Peyman Mestchian, M Makarov, and B Mirzai. 2005. Operational risk–COSO re-examined. *Journal of Risk Intelligence* 6, 3 (2005), 19–22.

[75] Edward A Morse and Vasant Raval. 2008. PCI DSS: Payment card industry data security standards in context. *Computer Law & Security Review* 24, 6 (2008), 540–554.

[76] Jason R.C. Nurse, Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. 2019. 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 (Cyber SA)*. IEEE.

[77] Kjartan Palsson, Steinn Gudmundsson, and Sachin Shetty. 2020. Analysis of the impact of cyber events for cyber insurance. *The Geneva Papers on Risk and Insurance-Issues and Practice* 45 (2020), 564–579.

[78] Jaeyoung Park, Woo-Jin Jung, and Beomsoo Kim. 2016. The Effect of Information Security Certification Announcement on the Market Value of Firms. *Journal of Information Technology Services* 15, 3 (2016), 51–69.

[79] Wolter Pieters, Sanne Hg Van Der Ven, and Christian W Probst. 2012. A move in the security measurement stalemate: Elo-style ratings to quantify vulnerability. In *Proceedings of the 2012 New Security Paradigms Workshop*. 1–14.

[80] Richard A. Posner. 2001. *Frontiers of Legal Theory*. Hardvard University Press.

[81] Roscoe Pound. 1968. *Social control through law*. Archon Books.

[82] Michael Power. 2007. *Organized uncertainty: Designing a world of risk management*. Oxford University Press on Demand.

[83] Sazzadur Rahaman, Gang Wang, and Danfeng Yao. 2019. Security Certification in Payment Card Industry: Testbeds, Measurements, and Recommendations. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 481–498.

[84] Eric Rescorla. 2005. Is finding security holes a good idea? *IEEE Security & Privacy* 3, 1 (2005), 14–19.

[85] Sasha Romanosky. 2016. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity* 2, 2 (2016), 121–135.

[86] Sasha Romanosky, David Hoffman, and Alessandro Acquisti. 2014. Empirical analysis of data breach litigation. *Journal of Empirical Legal Studies* 11, 1 (2014), 74–104.

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

[88] Jukka Ruohonen and Kalle Hjerppe. 2021. The GDPR enforcement fines at glance. *Information Systems* (2021), 101876.

[89] Bruce Schneier. 2014. Should US hackers fix cybersecurity holes or exploit them? *The Atlantic* (2014).

[90] Annika Selzer. 2021. Practitioners' Corner The Appropriateness of Technical and Organisational Measures under Article 32 GDPR. *European Data Protection Law Review* 7, 1 (2021). https://doi.org/10.21552/edpl/2021/1/16

[91] Annika Selzer, Daniel Woods, and Rainer Böhme. 2021. Practitioners' Corner An Economic Analysis of Appropriateness under Article 32 GDPR. *European Data Protection Law Review* 7, 3 (2021). https://doi.org/10.21552/edpl/2021/3/15

[92] Ravi Sen and Sharad Borle. 2015. Estimating the contextual risk of data breach: An empirical approach. *Journal of Management Information Systems* 32, 2 (2015), 314–341.

[93] Scott Shackelford, Anne Boustead, and Christos A Makridis. 2021. Defining "Reasonable" Cybersecurity: Lessons from the Public and Private Sectors. *Available at SSRN: https://ssrn.com/abstract=3919275* (2021).

[94] James Shires. 2018. Enacting expertise: Ritual and risk in cybersecurity. *Politics and Governance* 6, 2 (2018), 31–40.

[95] Mikko Siponen and Robert Willison. 2007. A critical assessment of IS security research between 1990-2004.

[96] Kyle Soska and Nicolas Christin. 2014. Automatically detecting vulnerable websites before they turn malicious. In *23rd USENIX Security*

Symposium (USENIX Security 14). 625–640.

[97] Jonathan M Spring, Tyler Moore, and David Pym. 2017. Practicing a science of security: a philosophy of science perspective. In *Proceedings of the 2017 New Security Paradigms Workshop*. 1–18.

[98] Kiran Sridhar, Allen Householder, Jonathan M. Spring, and Daniel W. Woods. 2021. Cybersecurity Information Sharing: Analysing an Email Corpus of Coordinated Vulnerability Disclosure. In *Workshop on the Economics of Information Security*.

[99] Kiran Sridhar and Ming Ng. 2021. Hacking for good: Leveraging HackerOne data to develop an economic model of Bug Bounties. *Journal of Cybersecurity* 7, 1 (2021), tyab007.

[100] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, you have a problem: On the feasibility of large-scale web vulnerability notification. In *25th USENIX Security Symposium (USENIX Security 16)*. 1015–1032.

[101] Samaneh Tajalizadehkhoob, Tom Van Goethem, Maciej Korczyński, Arman Noroozian, Rainer Böhme, Tyler Moore, Wouter Joosen, and Michel van Eeten. 2017. Herding vulnerable cats: a statistical approach to disentangle joint responsibility for web security in shared hosting. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 553–567.

[102] Leonie Maria Tanczer. 2016. Hacktivism and the male-only stereotype. *New Media & Society* 18, 8 (2016), 1599–1615.

[103] Colin Tankard. 2016. What the GDPR means for businesses. *Network Security* 2016, 6 (2016), 5–8.

[104] David Thaw. 2015. Data Breach (Regulatory) Effects. *Cardozo Law Review De-Novo* (2015), 151–164.

[105] Mark Van Hoecke and Mark Warrington. 1998. Legal cultures, legal paradigms and legal doctrine: towards a new model for comparative law. *The International and Comparative Law Quarterly* 47, 3 (1998), 495–536.

[106] Marie Vasek, John Wadleigh, and Tyler Moore. 2015. Hacking is not random: a case-control study of webserver-compromise risk. *IEEE Transactions on Dependable and Secure Computing* 13, 2 (2015), 206–219.

[107] Marie Vasek, Matthew Weeden, and Tyler Moore. 2016. Measuring the impact of sharing abuse data with web hosting providers. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*. 71–80.

[108] Vilhelm Verendel. 2009. Quantified security is a weak hypothesis: A critical survey of results and assumptions. In *Proceedings of the 2009 Workshop on New Security Paradigms (NSPW 2009)*. ACM, 37–50.

[109] Cesar G Victora, Jean-Pierre Habicht, and Jennifer Bryce. 2004. Evidence-based public health: moving beyond randomized trials. *American Journal of Public Health* 94, 3 (2004), 400–405.

[110] Robert K Vischer. [n.d.]. Legal advice as moral perspective. *Georgetown Journal of Legal Ethics* ([n. d.]).

[111] Tawei Wang, Karthik N Kannan, and Jackie Rees Ulmer. 2013. The association between the disclosure and the realization of information security risk factors. *Information Systems Research* 24, 2 (2013), 201–218.

[112] Jessica Weinkle. 2020. Experts, regulatory capture, and the "governor's dilemma": The politics of hurricane risk science and insurance. *Regulation & Governance* 14, 4 (2020), 637–652.

[113] Spencer Wheatley, Thomas Maillart, and Didier Sornette. 2016. The extreme risk of personal data breaches and the erosion of privacy. *The European Physical Journal B* 89, 1 (2016), 7.

[114] Norbert Wiener. 1954. *The human use of human beings: Cybernetics and society*. Number 320. De Capo Press.

[115] Dag Wiese Schartum. 2016. Making privacy by design operative. *International Journal of Law and Information Technology* 24, 2 (2016), 151–175.

[116] Josephine Wolff and Nicole Atallah. 2021. Early GDPR Penalties: Analysis of Implementation and Fines Through May 2020. *Journal of Information Policy* 11 (2021), 63–103.

[117] Josephine Wolff and William Lehr. 2018. Roles for Policy-Makers in Emerging Cyber Insurance Industry Partnerships. 46th Research Conference on Communication, Information and Internet Policy (TPRC 46).

[118] Daniel W Woods and Rainer Böhme. [n.d.]. Incident Response as a Lawyers' Service. *IEEE Security & Privacy* ([n. d.]), in print.

[119] Daniel W Woods and Rainer Böhme. 2021. How Cyber Insurance Shapes Incident Response: A Mixed Methods Study. In *Workshop on the Economics of Information Security*.

[120] Daniel W Woods and Rainer Böhme. 2021. SoK: Quantifying Cyber Risk. In *IEEE Symposium on Security and Privacy (SP)*. Oakland, CA, 211–228. https://doi.org/10.1109/SP40001.2021.00053

[121] Daniel W Woods and Tyler Moore. 2020. Does Insurance Have a Future in Governing Cybersecurity? *IEEE Security Privacy* 18, 1 (Jan 2020), 21–27. https://doi.org/10.1109/MSEC.2019.2935702

[122] Daniel W Woods and Jessica Weinkle. 2020. Insurance definitions of cyber war. *The Geneva Papers on Risk and Insurance-Issues and Practice* (2020), 1–18.

[123] Jack Wroldsen. 2015. Creative Destructive Legal Conflict Lawyers as Disruption Framers in Entrepreneurship. *U. Pa. J. Bus. L.* 18 (2015), 733–788.

[124] Jing Zhang, Zakir Durumeric, Michael Bailey, Mingyan Liu, and Manish Karir. 2014. On the Mismanagement and Maliciousness of Networks.. In *Network and Distributed System Security Symposium (NDSS 2014)*.

,