

How Cyber Insurance Shapes Incident Response: A Mixed Methods Study

Daniel W. Woods and Rainer Böhme

May 24, 2022

Abstract

Cyber insurance policies commonly indemnify the cost of incident response services. This creates a multi-layered economic problem in that the policyholder hiring external firms incurs *transaction costs* and the insurer paying the bill creates a *principal-agent problem*. We adopted a multi-stage research design to understand how insurers address the problem. First, we iteratively derived 12 stylised facts from 29 expert interviews and a sample of 480 partnerships with incident response firms made by 24 insurers. Second, we validated these facts via a workshop attended by 61 unique participants. The results show insurers have created a *private ordering* by controlling which firms are selected, negotiating prices ahead of time, and punishing low service quality by withholding future work. A minority of firms win the majority of work, thereby building trust through repeated interactions. We discuss how the findings relate to the economics of incident response, cyber insurance as governance, and ransomware.

1 Introduction

Cyber insurance allows firms to transfer cyber risk to an insurer. This creates a situation—known as a principal-agent problem—in which the agent (the policyholder) can make decisions that negatively impact the principal (the insurer). Early research predicted that insurers would address the problem by offering incentives for ex-ante security investments that reduce the likelihood of a claim [1–4]. So far, this has been undermined by an over supply of insurance and a lack of knowledge about which investments effectively reduce risk [5]. In actuality, the most significant intervention sees insurers indemnify the cost of incident response (IR) services [6–8].

Doing so opens up a Pandora’s box of economic problems. Why do insurers pay for external services and not offer subsidies for internal response? How do insurers ensure the policyholder selects an effective firm and negotiates a reasonable contract? Who is responsible for monitoring service quality?

DW Woods and R Böhme. *How Cyber Insurance Shapes Incident Response: A Mixed Methods Study*. The 20th Workshop on the Economics of Information Security (WEIS 2021).

These questions are naturally framed within transaction costs economics. Ronald Coase’s theory of the firm [9] speaks to when services are contracted on the market and when they are organised within the firm. Coase argues that firms emerge to avoid *transaction costs* associated with finding, negotiating, and monitoring service contracts tendered on the market. Coase’s theory suggests firms will hire external services if the associated transaction costs can be managed. Applying the same logic to cyber insurance requires care because the agent who transacts the service has different incentives to the principal who pays the bill.

This paper aims to build a concise and correct description of how cyber insurance solves the principal-agent problem. We adopt a multi-stage research design to describe market structure, processes and artefacts. The first stage uses unstructured expert interviews and insurers’ marketing materials to inductively derive 12 stylised facts. The second stage validates the facts via an online participatory workshop attended by 61 unique participants, which led us to reject one stylised fact. We classify our findings into Coase’s three types of transaction cost (search, negotiation, and monitoring [9]) and a fourth category related to market structure.

We discover that cyber insurance exerts significant control over the hiring of incident response firms by policyholders. Insurers control which firms receive work and condition this on: low hourly rates, harmonious interactions with other IR firms, and service quality that is sufficient to avoid disputes. The enforcement power of controlling who wins work and the associated system of (often unwritten) rules represents a *private ordering* [10]. The increased access to incident response services should be celebrated. However, the story should be qualified by a number of dysfunctionalities rooted in interpersonal relationships, market concentration, information asymmetries, and occupational licensing.

We identify relevant theory and prior work in Section 2. We describe our methodology in Section 3. The results are presented in Section 4. We discuss the implications and validity in Section 5, and then conclude in Section 6.

2 Related Work

The first decade of cyber insurance research predominantly introduced models of rational actors seeking to maximise continuously differentiable utility functions [11]. This approach followed luminaries of mainstream insurance economics like Kenneth Arrow [12] and Gary Becker [13]. These assumptions faced critiques from within mainstream economics field and later by empirical research into cyber insurance.

2.1 Transaction Cost Economics

We cannot recount all critiques of neoclassical assumptions, instead we focus on those that are most relevant to our phenomenon of interest, namely how incident response services are contracted. Oliver Williamson [14] argued that

optimal transacting between economic agents is limited by a bundle of concepts based around bounded rationality. Oliver Williamson [14] argued that although an optimal contract may exist in theory, it is hard to achieve in practice because of human limitations in “knowledge, foresight, skill, and time” [15, p. 199]. Achieving the optimal contract would incur time costs searching for and negotiating with counter-parties, and then over coming incomplete information to predict which possible complications require clauses and to then monitor adherence. The inability to do so is known as *bounded rationality*, and it leads firms *satisfice* [16] to accept sub-optimal contracts.

The resulting contractual defects may expose each party to strategic opportunism. A service provider may exploit contractual ambiguities to provide lower-quality service than the contracting party anticipated during negotiation. Sellers are exposed to *asset specificity*—the costs of re-deploying productive capacity for alternative uses. An opportunistic buyer could gradually reduce prices and exploit the seller’s cost of switching to alternative work [14].

These abstract ideas can be illustrated with incident response. Firms exhibit *bounded rationality* when responding to an incident because there is little time to collect information when an unknown adversary has exploited a system. *Incomplete contracting* is inevitable given neither party knows ex-ante the sophistication of the attack or the scope of the damage. An external incident response firm could engage in opportunism by over-billing hours or conducting a sub-standard investigation. Turning to *asset specificity*, a lawyer who specialises in advising on data privacy notification requirements incurs costs in switching to litigating on family law if cyber incident response work dries up.

Focusing on this collection of concepts—bounded rationality, incomplete contracting, strategic opportunism, and asset specificity—should make one pessimistic about organising production via transactions between firms. In contrast, Williamson [14, p. 141] introduces the concept of a *private ordering* to describe the ad-hoc solutions firms use to avoid and resolve disputes without the need for costly court proceedings. A private ordering provides a cost-effective alternative to contracts enforced by legal proceedings [17].

2.2 Empirical Cyber Insurance Work

Empirical work shows discrepancies with the neoclassical framing of insurers optimising production functions. Theoretical works frequently assume insurers can reliably condition cyber insurance prices or availability on the insured’s security level [3, 4, 18–22]. In actuality, insurers still use qualitative methods like questionnaires and telephone interviews to collect underwriting information [8, 23, 24] and are often limited by market conditions in doing so [5]. Pricing algorithms are crude and derived via methods like copying competitors and even guesswork [23]. The academic community should not be surprised given our own failure measuring security and linking it to risk outcome [25, 26].

Turning to post-incident, theoretical modelling suggests insurers should conduct forensic investigations to discover whether security levels were misreported in the application [27, 28] (a form of *strategic opportunism*). Dambra et al. [29]

cite an obscure blog to claim that “cyber insurance does not normally cover when employee errors are the cause of a malware infection”, which suggests insurers investigate claims with a view to denying losses. Yet one of the leading cyber insurance providers reports that most claims can “be traced back to a phishing email” [30] and another reports that 39% of claims result from employee error or social engineering¹. An analysis of over a hundred cyber insurance policies provides no evidence that employee errors are excluded [23]. Even when insurers try to “confirm or deny coverage”, they do so by asking questions by phone [8] because forensic investigations are costly (a form of *bounded rationality*).

Insurers are, however, willing to indemnify the cost of incident response firms [6–8] to act in the policyholder’s interest, which means information collected is not used to deny coverage. Such firms represents the majority of the cost of 70 cyber insurance claims [31]. Wolff and Lehr [32] discover that insurers offer more partnerships with legal firms than technical.

Although recent empirical work discovers many quirks of the cyber insurance market that diverge from theoretical assumptions, a collection of anecdotes is not knowledge. The next section describes a research design that aims to distil general insights.

3 Methods

To obtain a concise and correct description of the market, we first collected information via published documents (Section 3.1) and interviews (Section 3.2). We used the resulting data to derive a set of 12 *stylised facts*. Stylised facts are essentially true but fail to explain certain particulars. They can be seen as a starting point for theory construction in economics [33]. We then organised an online workshop and encouraged participants to contradict our stylised facts (Section 3.3).

In pursuit of market realism, the first stage collects data directly from market participants and the second stage uses market participants’ feedback as a falsification criterion. This assumes: (i) the participants share a common view of reality, and (ii) we can reliably interpret participants’ reports. Reliable interpretation is achievable because our object of enquiry is market processes and artefacts that are experienced and discussed often using specific terminology. If the first assumption (i) does not hold, our stylised facts would be rejected when presented to more market participants in the validation workshop.

Section 3.1 describes how we manually extracted information about corporate relationships from the websites and marketing materials of cyber insurance carriers. Section 3.2 describes our approach to recruiting, conducting and analysing expert interviews. Section 3.3 describes the organisation of an online validation workshop. Throughout we adopt key terms used by participants to improve fidelity.

¹<https://chubbcyberindex.com/#/incident-growth>

Key Terms

- **Incident response firm:** Any external firm engaged after the client suspects a cyber incident has occurred. Services provided include legal advice, forensic investigations, IT recovery services, credit monitoring, notification logistics, public relations advice, and forensic accounting.
- **Panel:** The list of firms the insurer has partnered with.
- **Hot line:** The dedicated phone line that policyholders are instructed to call upon discovering a cyber incident.
- **External counsel:** An external law firm hired to provide advice to and/or represent a victim firm.
- **Discovery:** The legal mechanisms resulting in compulsory disclosure, at a party's request, of information that relates to the litigation [34].
- **Client-attorney privilege:** The client's right to refuse to disclose and to prevent any other person from disclosing confidential communications between the client and the attorney [34].

3.1 Public Relationships

Prior work [32] shows that many insurers advertise a list of incident response providers covered by the cyber insurance policy. These providers are variously described as *preferred*, *pre-approved* or *authorized* providers, and a *global partner network* depending on the insurer. Going forward we will use *insurer's panel* to describe all providers who the policyholder needs no prior approval to use. In addition, most insurers say alternative providers may be used with prior-approval on a case-by-case basis.

We used lists² of US cyber insurance carriers as a seed sample and a search engine to find each insurer's preferred providers. This involved searching the insurer's sub sites and extracting documents, such as brochures or policy wordings, describing the cyber insurance products. We captured the panels of 24 insurers advertising 480 preferred providers of which 151 were unique. No automated data extraction was conducted.

3.2 Expert Interviews

We then conducted interviews to collect contextual information.

Recruitment We initially recruited participants through our networks, which resulted in a handful of participants and a notable lack of technical vendors. Rather than begin cold emails, we shared an advert on [linkedin.com](https://www.linkedin.com) asking

²For example: <https://www.reinsurancene.ws/top-20-us-cyber-insurance-companies/>

for potential participants to get in touch, by which we recruited 19 participants. The remaining participants were recruited via snowball sampling.

The advert, which can be found in the Appendix, shared a preliminary figure based on the data from Section 3.1. The rest of the post explained the aims of the study, the funding source, and relevant hashtags. LinkedIn’s engagement statistics report 5K views, 13 comments, and 15 reshares.

This mix of recruitment channels may introduce biases. At the individual level, recruiting via LinkedIn biased our sample towards professionals who value and participate in online networking. To address incomplete coverage of firms, we targeted our recruitment towards firms we had not yet spoke to in the top US cyber insurance carriers and the service providers listed on their website. We stopped recruiting when further interviewers resulted in similar reports to what had already been collected. We spoke to 10 insurance professionals, 13 IT practitioners, 5 lawyers and 1 recruiter in this ecosystem.

Interview Procedure Interviews were scheduled to last 60 minutes and were conducted by video call. We made notes during the interview and also recorded the video and audio if the participant provided written consent. We obtained ethical approval for the interviews from our institution, which included reviewing the information sheet, consent form, and interview script.

The interview guidelines, which can be found in Appendix B, were drafted after pre-study discussions with a range of stakeholders. We followed two separate documents for service providers and individuals who selected service providers (e.g. insurers and external counsel). In addition, we asked follow-on questions and also for corroboration of statements made by other participants. The scripts contained a general section to understand the participant’s background and one section for each of the three types of transaction cost (search, negotiation, and enforcement). A few interviews were conducted with other actors (e.g. brokers, recruiters, re-insurers) and we modified the scripts to make the question relevant.

Some individuals wanted to contribute to the study but were uncomfortable with signing a consent form and/or being a recorded. We provided additional anonymity to such participants by not audio recording the interview or quoting anything from the notes. This research data was treated the same in terms of data protection, revocation and deletion.

Analysis We followed an iterative process, which involved: conducting interviews; writing up a description of the market that explained previous discussions; and, presenting aspects of this description to participants in subsequent interviews. We converged on a tolerable level of generalisation. For example, all insurers draft a list of approved providers (a panel) but we failed to build a general account of the process by which firms were added to panels. This means our findings were exposed to unstructured falsification even before the validation workshop, which we now describe.

3.3 Validation

Given that the previous research methods were primarily inductive, we designed an online workshop to validate the interview findings. This involved presenting the findings to practitioners and specifically requesting refutation. We advertised the workshop via professional networks (LinkedIn and Twitter).

The online workshop consisted of a video stream of the researchers superimposed onto slides with live annotations. The chosen platform (Twitch) displays a live chat alongside the stream, and the latency was sufficient for the live presentation to incorporate chat comments. Choosing a relatively niche (gaming) platform allowed the participants to choose pseudonyms to protect their anonymity, and also created the potential for serendipitous participation from Twitch users with no prior interest in the study. Two chat moderators were online but did not have to intervene. The recording was made available afterwards for asynchronous viewing and also as a research artefact³.

The stream began with introductory slides and then broke our findings down into four sections (search, negotiation, monitoring, and market structure). For each theme, we presented slides explaining the stylised facts in more detail and then paused for 60 seconds, specifically asking the audience to comment on the stylised facts. The recording was available in the following days and the audience could feedback via a survey instrument or directly contacting the author.

The platform reports participation statistics including: 61 unique viewers, 17 unique chatters, and 96 messages in chat. The majority of viewers watched from the US (40%) and the UK (15%) with 22% coming from the researchers' country, all of whom we assume to have been members of our research group. The recording was viewed an additional 65 times. Very few viewers opted to fill out the survey we prepared.

After the recording, we collected all messages from the chat and discarded any from our research team. The remaining messages were classified into pleasantries and jokes, questions, points of information, positive confirmations, and refutations. We avoided two failure modes; (i) no refutation at all (a sign findings are not clear enough to contradict), and (ii) constant refutation. The feedback led us to reject and replace one stylised fact, to change the emphasis around another issue, and also to increase confidence in other findings via explicit confirmation.

There are question marks over this stage of our study. The lack of gate-keeping means we do not know the participants' experience or affiliation, nor could we ensure non-intersection between the sets of individuals who generated and validated the facts. The online format made it difficult to sense whether participants comprehended the facts or whether the participants influenced each other. Only time and future research will tell. However, some benefits of this innovation are clear: (i) the existence of an independently verifiable research artefact; (ii) avoiding the elitism by which researchers arbitrarily divide the world into experts and non-experts; (iii) eliminating disease transmission risk; and (iv) offering the final word to the community that we aim to describe.

³<https://www.twitch.tv/videos/908724413>

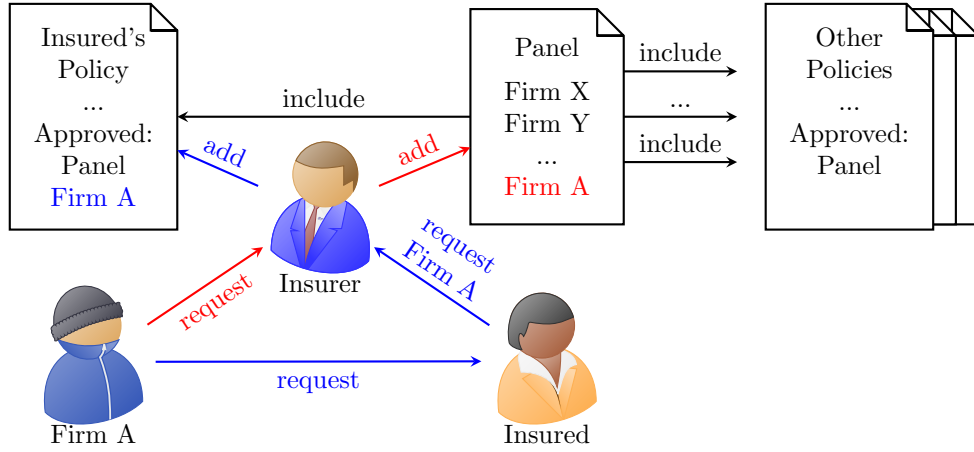


Figure 1: Firms gain prior approval by either being added to one insured’s policy (blue) or to the panel (red). The panel applies to many policies.

3.4 Ethics

Our biggest ethical concern was inadvertently damaging the careers of participants. With this in mind, we anonymised participants’ names, job roles and firms and also avoided certain topics. For example, we never asked questions about individuals and moved onto a different topic when one participant became hesitant when talking about ransom procedures. We did not anticipate risk of psychological damage talking about abstract market procedures.

Our secondary ethical concern was wasting participants’ time given we offered no concrete rewards (e.g cash, amazon vouchers, or raffle entry) for participation. As a result, we saw our duty as conducting high-quality research and sharing it widely. The validation workshop was also designed to disseminate information.

4 Results

This section describes how insurers address the problems of search, negotiation and monitoring, and then moves onto how this influences the market structure of firms providing incident response services.




Figure 2: Size of each insurer’s panel for 14 of the top 20 US cyber insurance carriers who make it publicly accessible (in descending order). The Appendix contains the same table with all insurers in our sample.

Search

- S:1 Insurers build a panel of firms whose services the policy will indemnify, and the hot line operator triages by recommending specific providers.
- S:2 Shortlisting for the panel is selective and the provider must commit to certain terms (e.g. hourly rate or fixed price for certain investigations).
- S:3 Most firms follow the recommendation of the hot line operator, who tends to be an external law firm in the US.

Search Insurers structure policies and processes to influence who is selected because IR providers vary in both quality and cost. Stylised fact *S:1* describes the two-step selection process by which insurers build a panel of firms whose services the policy will indemnify, and then ask insureds to contact a hot line allowing the operator to recommend a firm from the panel.

Figure 1 shows the two routes for IR firms to be affirmatively covered by a cyber insurance policy. Firms may apply to join the *panel*, which is included in every policy. Alternatively, the insured may request a specific firm, which may occur if the IR firm has been used in the past or when the IR firm sells/manages a product in the insured’s network.

Stylised fact *S:2* says the process of joining a panel is selective and the provider must commit to certain terms (e.g. hourly rate or fixed price for certain investigations), which can be a lengthy process. One IR firm reported exchanging documents for over a year before giving up on working with that




Figure 3: The hot line operator recommends firms from the panel, and some law firms even hire the firms under the direction of counsel.

insurer. In contrast, adding a firm to the policy only required informal negotiations about price and experience. The arrows in Figure 1 flip direction when insurers expand into new markets and actively search for partners. Two providers reported that after becoming established in the insurance ecosystem, the insurer made first contact in 20 – 30% of the panels they were added to. Other insurers had to actively search for partners outside the US/UK.

Cyber insurers often publicly advertise panels of providers as a marketing tool. Figure 2 shows that 70% of the Top 20 US cyber insurers⁴ advertise their panels publicly. Larger insurers tend to list more firms because they have more work to distribute. We further classify the firms into the services offered (legal, forensics, and communications), which we analyse in more detail when we turn to market structure.

Most of the panels in Figure 12 contain multiple providers for each category, which raises the question of how a single provider is selected. Insurers control this decision via a *hot line* operated by either the insurer or a third party as depicted in Figure 3. Third parties may be appointed for reasons including logistics (24/7 multi-lingual call centres), general experience dealing with cyber incidents, or occupational license. Stylised fact *S:3* holds that most firms follow the recommendation of the hot line operator, which tends to be an external law firm in the US.

Ensuring insureds hire the recommended firm is not a trivial outcome. “Bait-and-switch” was evocatively used to describe how most insureds end up with a cheaper, less reputable firm even though the panel contains famous forensic firms. Less cynically, different firms excel in specific incidents and the hot line operator functions to match incidents to the right provider. Multiple forces push insureds towards the recommended firm: superstar firms have limited capacity

⁴<https://www.reinsurancene.ws/top-20-us-cyber-insurance-companies/>




Figure 4: Hot line operators oversee negotiations, which rely on ex-ante agreements and contract templates. The result is then approved by the insurer.

and avoid working on vanilla incidents, insureds must pay fees up-front and face cash flows problems, on-panel firms are instructed to direct insureds back to the insurer if contact was independently made, and some insurers do not publicly list famous firms even though they are used for complex incidents.

Negotiation

- N:1 Insurers negotiate hourly rate/fixed pricing while building the panel, policyholders provide information about their environment (e.g. number of sites or machines), and hot line operators advise on the scope of work. This results in a statement of work, which must be approved by the insurer or a delegated authority.
- N:2 Often insureds directly contract with external counsel, who then hire firms on the insured's behalf. Technical work may be further sub-contracted, especially for high risk activities like ransomware negotiation and payment.
- N:3 Insureds negotiate additional services that are not covered by cyber insurance. For example, monitoring tools installed as part of the investigation are often retained by the insured at their own cost.

Negotiation Given the scope for expensive or unnecessary services, insurers also exert influence over negotiations as depicted in Figure 4 and in stylised fact *N:1*. The cost of investigations, such as hourly rate or the total price for simpler investigations, is negotiated as the insurer builds the panel. Hot line operators




Figure 5: Insurers track day-to-day progress on the investigation via external counsel’s summary of the vendor’s informal reports. The report is not always written, and it is even less frequently sent to the insurer.

and forensic firms work together sufficiently often that contract template can be used. This means insureds only need to provide information about the incident and environment to be investigated (e.g. number of sites or machines). The resulting work contract must then be approved by the insurer, although this authority is sometimes delegated to the hot line operator.

Stylised fact *N:2* raises the possibility that services may be further subcontracted. *Often insureds directly contract with external counsel, who then hire firms on the insured’s behalf (N:2)*. This practice is a result of US law. Litigants suing the victim firm can use *discovery* to obtain documents that are relevant to the case, such as the forensics report. Law firms argue that by hiring forensics firms, it is easier to argue the report was produced in anticipation of litigation and so it is protected by *client-attorney privilege*.

Risk associated with negotiating and paying ransoms provides a second rationale for subcontracting. One participant explained a regular arrangement in which one firm investigated ransomware incidents, out-sourced the negotiation to another firm, and then a third firm facilitated the payment.

Stylised fact *N:3* shows that IR firms may have an ongoing influence on the victim’s security posture by up-selling mitigation products. This creates an unanticipated sales channel with conversion rates upwards of 50% (reported by three separate participants). One IR firm originally bought subscriptions to end-point products used as part of investigation at open-market prices. Upon realising that most insureds pay a subscription fee to keep the product after the investigation, the IR firm negotiated a deal with the product company in which the IR firm kept all of the on-going subscription above a fixed price.

Monitoring

- M:1 Preferred providers self-monitor in order to avoid disputes and receive future work from insurers and external counsel. There are few disputes when on-panel firms are used, but going off panel frequently results in incident response services not being indemnified.
- M:2 Insurers rely on external counsel to monitor providers on a day-to-day basis. Further, the insurer mainly receives informal/verbal reports to avoid documents that could be discovered by a litigant.
- M:3 Forensics reports are not standardised, and so investigations are structured according to the law firm/lawyer.

Monitoring Once a provider has been selected and a contract negotiated, service provision should be monitored. Insurers tend not to micro-monitor each contract, instead the decision to award work to providers is linked to past performance across multiple claims. Stylised fact *M:1* shows this system is broadly successful. Many participants who regularly won work via this ecosystem struggled to respond to the question “*What kind of disputes arise between insurer and service provider?*” Participants working in claims departments reported that the majority of disputes resulted from insureds hiring off-panel firms. Such disputes most commonly concerned the final bill. It is unsurprising on-panel firms avoid such disputes given prices are negotiated ahead of time (*S:2*).

Stylised fact *M:2* shows that the performance of forensics firms is primarily monitored by external counsel on a day-to-day basis (see Figure 5). Further, *the insurer mainly receives informal/verbal reports to avoid documents that could be discovered by a litigant (M:2)*. These findings cast doubt over the ability of insurers to link forensics performance to assigned work. In actuality, a legal professional’s unstructured evaluation functions as a proxy for the quality of investigations. Such 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). This cynicism should be qualified given such lawyers regularly work with forensics providers and some have even pursued formal information security training.

Finally, *the investigation’s deliverable—the forensics report—is not standardised, and so investigations are structured according to the desires of the law firm/lawyer (M:3)*. One participant reported that forensics firms share spreadsheets outlining how law firms and even individual lawyers want investigations to be presented. While this structure no doubt improves how efficiently external counsel can advise on regulatory and litigation risk, it is less clear whether it is appropriate for addressing and learning from technical risk. This problem is compounded when client-attorney privilege means the investigations is not formally documented. One insurer reported receiving a forensics report in less than 10% of their claims.




Figure 6: The distribution of number of listings per provider broken down by category.




Figure 7: Law firms with more than two listings.

Market Structure

C:1 A handful of law firms dominate. A larger number of forensics firms receive work, such firms tend to be service rather than product based.

C:2 Technical providers are often replaced mid-way through an investigation.

C:3 There is always upstart forensics firms offering a lower price. Often such firms are founded/led by the former employees of dominant firms.

Market Structure We now provide observations about which firms win work and how this changes over time. Figure 2 showed the number of firms on each insurer's panels. Figure 6 flips the analysis and shows a handful of providers receive the majority of listings, which led us to derive stylised fact *C:1*

Figure 7 shows the distribution of listings among lawyers. Mullen Coughlin,




Figure 8: Forensics firms with more than two listings.

who only take data privacy cases, were listed by 80% of the insurers in our sample even though they were founded in 2016. The four legal firms with the most public listings in our sample collectively managed over 3 500 incidents in 2018 [35]. All four describe themselves as breach coaches. The trademark is owned by a firm in the ecosystem (NetDilligence).

Firms with forensics capabilities hold the majority of the listings (see Figure 6 and Figure 8) but there is less concentration when compared to law firms. Forensics vendors follow surprisingly diverse strategies. Charles River Associates and Kroll were founded in 1947 and 1972 respectively, whereas Ankura and Crypsis were founded in 2014 and 2015. Product-based IR firms (e.g. FireEye and CrowdStrike) build or own products that are relevant to the investigation. Service-based IR firms who have no pre-incident access to the victim’s environment were more common and received more work (*C:1*).

The communications category (Figure 9) should be divided into strategic and logistical services. Public relations (PR) consultants like Edelman or FleishmanHillard provide advice on how to communicate with the public at large, whereas credit monitoring and logistics firms like Epiq and Experian provide direct notification and services to individuals. PR services are only required when the media is likely to cover an incident, which tends to be for the largest firms. However, notification is a regulatory requirement following a breach and so many more firms require these services. Similar to the distinction between service and product-based IR firms, notifying individuals relies on infrastructure (e.g notifying millions of breach victims by mail) and credit monitoring requires access to a credit bureau, whereas PR services only require consultants and their




Figure 9: Communications firms with more than two listings.

professional networks. Although each of the only three major credit bureaus in the US (Experian, TransUnion and Equifax) were listed by insurers, Experian were the only one to be listed by a majority (80%) of insurers. It is worth noting we did not interview anyone from this category of firms.

In the validation workshop, multiple participants contradicted the stylised fact *C:2* “technical providers are often replaced mid-way through an investigation”:

30:30 craifdmb4ever: I think replacement of investigators is relatively rare

32:17 adhontwitch: I also agree that its a very rare occurrence that someone gets replaced...

which led us to reject it. Participants suggested that this happens rarely, and firms are punished by not receiving work in the future when this happens more than twice. This supports interview reports that insurers would not award work to off-panel forensics providers who failed to impress in trial investigations. We did not hear about any failures on the part of law firms, most likely because identifying reporting requirements is a more certain task than identifying and containing an active adversary.

Finally, stylised fact *C:3* suggests IR firms struggle to maintain dominance. Figure 10 shows how many of the newer forensics providers are founded or run by employees of formerly dominant firms, whereas no communications firms and just one law firm (Mullen Coughlin) did so. Workshop participants explained junior moves are even more revealing with a claim that one forensics firm “lost 46% of talent to competitors with 17% of them going to Arete as an example”. Exhaustively tracking these dynamics would require a different methodology.

security practices in organisations. Section 5.3 considers the impact of cyber insurance on the ransomware epidemic. Finally, Section 5.4 reflects on the validity of our results and which aspects of our methodology should be used/avoided going forward.

5.1 Economics of Incident Response

It is useful to contrast two perspectives and the economic intuitions they invoke. *The democratisation of incident response* argues that cyber insurance carriers identify effective IR firms and use market power to drive down the price for policyholders, many of whom would not otherwise contract these services. In contrast, *the commoditisation of incident response* holds that downward pressure on prices has driven quality out of the market as forensics firms economise via automation and shallow investigations.

The market for lemons provides a simplistic explanation for the commoditisation position (as it has for other security economics problems [36–40]). Suppose a low quality investigation costs \$250 per hour while a high quality investigation costs \$500 per hour, and the insurer struggles to distinguish between the two (as stylised fact *M:2* suggests). Regardless of what hourly rate is negotiated, a rational investigator in a single-shot game conducts a low quality investigation and pockets the difference. Thus a rational insurer should pay \$250 per hour.

This logic would also apply to individual firms without insurance unless it can be argued they can better monitor IR service quality than insurers. Potential reasons could be found in the firm’s proximity to the investigation or in the reality that many large organisations possess more IT expertise than insurance companies, but neither explanation is particularly convincing.

It seems easier to argue that insurers have a comparative advantage in contracting incident response services. In the absence of insurance, a victim could be exposed to a hold-up problem [14] because the victim’s time investment in searching for and negotiating with a potential provider represents a sunk cost. This relationship-specific investments is much larger for the victim firm than an insurer because time spent searching can be used by the threat actor to cause more damage. Forensics firm could use this sunk cost to *hold-up* the victim for a higher price. Further opportunism might involve skimping on service quality or conducting an unnecessarily broad investigation. In contrast, insurers incur search costs once for all policyholders and negotiate before an incident has occurred to shift bargaining power (see stylised fact *S:1*). Additionally, insurers hiring firms can be seen as a repeated game in which trust relationships are developed across multiple claims that prevent post-contract hold-up issues (see stylised fact *M:1*).

A stronger argument for commoditisation focuses on investigation infrastructure and planning. NIST-800-61 [41] recommends that incident response integrates planning, monitoring and investigation in order to use data collected before the incident, but this requires data retention to be put in place ex-ante. Product-based IR firms can do this and insurers reported more efficient inves-

tigations, admittedly at higher hourly rates⁵. In contrast, the service based IR firms favoured by insurers must rely on whatever data collection and retention processes were in place. As a result, such services can be provided to any organisation—a sign of commoditisation.

This framing assumes that commoditisation is necessarily undesirable. Perhaps it is the right response to the commoditisation of cybercrime [44–46], especially given service-based IR can rely on the monoculture of corporate software. For example, most business email compromise involves an Office365 inbox, which helps explain why some firms can offer fixed price investigations. Corporate software providers could reflect on whether they do enough to support incident response. For example, one IR firm released an open-source tool⁶ based on undocumented APIs to investigate Office365 account activity. Microsoft subsequently restricted access to the API and all IR firms using that functionality had to develop new and less efficient methods.

The views can be reconciled by casting the insurer’s hot line as a form of triage that functions to match incidents with response firms. Commoditised incidents like Office365 account compromises or unsophisticated ransomware strains affecting small businesses may be more efficiently investigated via automated scripts, whereas multi-national corporations compromised by nation state actors likely require costly procedures like bespoke malware analysis and the data processing capabilities of product based firms. Some of this happens naturally—organisations targeted by nation state actors are likely to request a firm they have worked with in the past (see the blue line in Figure 1). Although a self-interested insurer would hire a sufficiently qualified firm, the lemons problem likely pushes insurers to favour cheaper firms at the margin. A very academic recommendation (read: easier said than done) is to develop metrics to track investigation outcomes and use these to evaluate triage decisions over time.

5.2 Cyber Insurance as Governance

Over the last twenty years, multiple authors and institutions have considered how *insurance as governance*—the idea that buying insurance changes how the policyholder manages risk—might apply to cybersecurity. This began with theoretical papers at the Workshop on the Economics of Information Security and soon began appearing in policy discussions [47]. Since 2012, the EU’s cybersecurity agency [48], the US Department of Homeland Security [49–51], the UK Government [52], the OECD [53] and more recently the US Senate [54] have discussed how to support cyber insurance as governance. Proposed policy measures include governments providing funds for insurers who have suffered catastrophic cyber losses—the US Treasury already provides such a back-stop for events of cyber terrorism [55] (whatever such events look like)—and making cyber insurance mandatory for SMEs [56].

⁵Modelling this dynamic as a lock-in problem [42, 43] seems like a promising direction for future work.

⁶<https://github.com/CrowdStrike/Forensics/tree/master/O365-Outlook-Activities>

As Section 2 described, there is little evidence cyber insurance rewards security practices [5, 57, 58]. The literature broadly agrees that cyber insurance as governance is most influential when it comes to ex-post response [6–8]. Talesh [7] provides the most complete picture, which is broadly in line with the democratisation of incident response from Section 5.1. Our findings both support and qualify aspects of this, as well as providing entirely new considerations.

Merely stating cyber insurers provide access to incident response services is an under-statement [5, p. 2]. More specifically, insurers govern the relationship between policyholders and incident response service firms. This involves: searching for and negotiating with providers ahead of time to gain discounts relative to open market rates (30% is typical based on our sample); concentrating work among a handful of firms enabling streamlined processes to emerge; and withdrawing future work from providers who do not deliver an expected quality of service. Even the disenfranchised IR firms would admit that cyber insurance *is* a form of governance over incident response services.

Dissenting voices would instead contend that insurers are not fair nor even effective governors. The processes by which IR firms are added to panels and recommended by hot lines are opaque at best and nepotistic at worst; insurer panels consist of firms known by the insurer because interpersonal trust built via conference or business interactions is perceived to be more reliable than impartial alternatives, such as certification⁷. Other points of contention include: the same accusations of nepotism levelled at external counsel, unsustainable hourly rates, narrow work contracts preventing in-depth investigations, and unwillingness to pay for remediation of security issues.

It is unsurprising insurers are unwilling to pay for remediation given insureds are willing to pay for security products and services not covered by the insurance contract (see stylised fact *N:3*). Firms may even be willing to incur a strategic loss on an investigation because interacting with victim firms opens up a sales channel. The lawyers we spoke to felt this was unprofessional. Highlighting disciplinary differences, the security community sees incidents as a valuable learning experience [60] and the professional duty is to improve the client’s security posture wherever possible. Thus, the choice of ex-post response provides indirect influence over future ex-ante mitigation.

Turning to the role of lawyers, Talesh argues cyber insurance processes are “less about simply avoiding being sued” and more about preventing incidents occurring. The comparison to employment practices liability insurance [7] could well hold, but the community should not under-estimate how legal risk shapes and even prevents ex-ante mitigation. Insurers appoint law firms at the top of the IR hierarchy (see stylised fact *S:3* and Figure 3) and considerations around client-attorney privilege prevent the documentation and sharing of forensics investigations (see stylised fact *M:3*). Quantifying the opportunity cost of squandered knowledge is impossible, but legal risk is no doubt limiting the ability of insurers to build knowledge over time.

⁷One could imagine an insurance contract offering to indemnify any provider certified to be at the right level. Certifying investigators may escape the documented problems certifying websites or software [37, 59].

Beyond simply complying with laws, this ecosystem resolves ambiguities in law on a daily basis. For example, victims must notify affected individuals following a business email compromises. This problem has a technical component (which emails were accessed?) and a legal component (does an accessed email contain personal data?). Automated solutions classifying whether the email contained personal data are often deployed to reduce the cost of manual analysis by a team of paralegals, which can cost up to \$500k. The inevitability of false negatives in this classification task sees the data subjects’ right to notification traded off against the cost of investigation, and yet the alternative of manually filtering every compromised inbox is hardly any better. Industry norms are being established one incident at a time, and the process is intensified by the concentration of work among a few firms with common solutions.

This supports the *new legal realists*’ view that ambiguities in law are often resolved outside court rooms in ways that can only be probed via empirical social science research [61]. The same methodological requirement seems to apply to understanding cyber insurance as governance. Our findings like investigations becoming a sales channel or client-attorney privilege distorting information flows were not observable via insurance policy analysis (e.g. prior work has analysed 3 [3], 6 [62], 14 [63] and 100+ [23] policies) nor were the findings modelled (let alone predicted) by theory.

5.3 Ransomware

Recent years⁸ have seen ransomware become one of the most economically significant cyber crimes. Prior research has focused on the design [66–69] and strategic decisions [70, 71] related to preventative measures. An alternative approach is to use public policy to change negotiation strategies [72]. The most relevant proposal to this paper is banning insurers from indemnifying ransom payments [73, 74].

This is not the first historical example in which insurance was accused of incentivising crime [75]. The industry managed fears that child life insurance would increase the prevalence of murder by emphasising the value in terms of support for grieving families [76] or by shunning “economic terminology” in favour of religious symbolism to placate nineteenth century moral intuitions [77]. Detractors would respond that this is about economics not moral intuitions.

The underlying logic of critics [78] holds that paying ransoms increase the likelihood of future ransoms as payments demonstrate profitability leading new actors to enter, or increases the impact as existing criminals increase ransom demands (known as ransom inflation). In the language of economics, the victim paying the ransom demand imposes a negative externality on peers who now face a higher threat level. At the margin, victims are more likely to pay if insurers indemnify some or all of the payment [79].

This reductive logic misses other impacts. For example, insurers concentrate negotiation and payment among a handful of firms ($N:2$), with one firm

⁸Ransomware’s academic heritage is much longer [64, 65].

(Coveware) reporting working 150 ransomware cases per month [80]. Market concentration in physical kidnap insurance enables the negotiation standards necessary to prevent ransom inflation [81]. Remarkably this arrangement of private actors maintains kidnap victim recovery rates of 97%+ without causing ransom inflation [81]. In comparison, Coveware report recovering the decryption key following payment in 99% of cases [82], but also report worrying ransom inflation with the mean payment going from \$6K to \$155K between Q3 2018 and Q3 2020 [83].

Coveware’s founder attributes the key recovery rate to being able to track ransomware actors across 1000+ yearly negotiations and to punish gangs who renege on agreements [80]. In this way, incident response firms share information across the firms they work with, which contributes to a long held public policy goal [84]. Perversely, the resulting trust created by repeat interactions may even be supporting the ransomware business model by making contracts enforceable, mitigating the disruptive potential of dishonour among thieves [38].

Thus, insurers concentrating response among a few firms could either improve negotiation discipline *or* increase trust in the criminal business model. More generally, whether insurers are worsening the epidemic is an empirical question and we are not aware of any such answers. This points to the wider problem of uncovering causality in insurance markets where selection effects are the name of the game.

A more modest goal is accurately describing what happens in the market, but even this is difficult. We failed to distil any stylised facts about ransom procedures as this varies so much across insurers and providers. For example, one insurer conducts sanctions risk assessments via block-chain analysis for insureds, presumably making payments *more* lawful. Another participant offered insureds interest-free loans for the ransom, which increases propensity to pay. Similarly, some of the forensics firms in Figure 8 pay ransoms, others will negotiate but not facilitate payment, and others will not even negotiate. This motivates considering the validity of our results.

5.4 Validity and Limitations

The validity of an exploratory study is hard to probe. It is clear that our analysis of insurer panels missed some firms. For example, SpearTip report working over a thousand insurance claims [35] but were listed by just one of the panels in our sample. Missing some firms is natural given our sample of 24 firms was not exhaustive, these lists are updated infrequently, and some insurers regularly send work to off-panel firms.

A more systematic bias is our inability to observe panels managed in the eRiskHub, which was used by many of the insurers who did not advertise a vendor list. In this multi-sided platform, the insurer controls which vendors and lawyers are shown to the insured. We were provided guest access to this platform and it lists many of the same providers as the panels in our sample, although we could not tell which firms were selected by which insurer. This

platform presents an interesting example of market power for future work⁹.

Turning to the data collected via expert interview, our aim was to generalise beyond anecdotes. Adopting the convention of stylised facts—statements that essentially true but fail to explain certain particulars—from economics allowed us to distil general findings without over-claiming generality. The validation workshop exposed the stylised facts to falsification. The research artefact speaks to the level of participation, the chat falsified only one fact, and we also circulated a version of this paper among participants. Thus, we have reasonable confidence that the stylised facts hold true in most cases.

Looking forward, online validation workshops are applicable whenever research findings speak to market experiences and structure that practitioners are better placed to observe, which we term *market realism*. Recruitment is difficult because highly specialised/paid professionals cannot be recruited on campus or via mechanical turk. Our experience suggests participants value independent analysis, which academics are well-placed to provide. The second issue to overcome is encouraging active participation. We did so with 60 second pauses that actively encourage comments, adding our own comments to the chat, and by responding to the chat throughout the live stream.

6 Conclusion

Cyber insurance providers go beyond providing access to incident response services. Insurers control who is hired by limiting coverage to on-panel firms and then directing policyholders to call a hot line upon discovering an incident ($S:1$). This market power is used to enforce a private ordering with multiple effects: (i) insurers can drive down hourly rates by linking negotiation to the volume of awarded work ($S:2$), (ii) contract templates emerge as providers work similar incidents repeatedly and processes become streamlined ($N:1$), and (iii) IR firms self-monitor service quality to avoid disputes and ensure future work ($M:1$). Together these effects function to reduce transaction costs and make externally provided incident response more economically efficient, which is especially beneficial for firms with low security maturity. We term this narrative the *democratisation of incident response*. This is analogous to insurers funding fire brigades before they existed as a public service [85].

This narrative is complicated by other distorting effects of the market power of insurers: (iv) insurers concentrate work among a handful of law firms and a small number of forensics firms ($C:1$), (v) law firms commonly lead the incident response and choose which firms to hire ($S:3$), (vi) the progress and results of investigations are inconsistently and informally reported to insurers to protect client-attorney privilege ($M:2-3$), and (vii) forensics firms may even be running a loss on investigations to open up a sales channel with clients ($N:3$). Evaluating these effects is more difficult. Insurers undoubtedly face incentives to concentrate work among the *most* efficient firms but may lack the information to do

⁹The platform is run by the same firm that holds the trademark under which the top four law firms in Figure 7 operate

so given insurers do not monitor service quality and must rely on second hand reports. The only thing insurers reliably monitor is cost, which may lead to the *commoditisation of incident response* that rewards cost-cutting above quality.

Perhaps most curiously, the US cyber insurance market tends to see finance professionals delegate incident response leadership to external counsel. As a result, legal professionals coordinate a team with expertise in digital forensics, public relations, and the logistics of notification and credit monitoring. The situation emerges because American courts extend a special power, namely client-attorney privilege, to individuals with the right occupational license. Protecting privilege was undoubtedly valuable when claims were driven by customer and shareholder lawsuits following a data breach, but it has less value during a ransomware epidemic when losses are driven by technical compromise. At present, insurers pay for a forensic investigation into what caused each compromise and then squander the opportunity to build a structured database that would enable analytical work in the future.

Acknowledgements

First and foremost, we thank the participants for volunteering their time to help us understand how the ecosystem works. We also thank the four WEIS reviewers, Shauhin Talesh, and Jono Spring for detailed and insightful comments. We also received useful feedback from the Cyber Insurance Special Interest Group at FIRST. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 894799.

References

- [1] Bruce Schneier. Insurance and the computer industry. *Communications of the ACM*, 44(3):114–114, 2001.
- [2] Lawrence A Gordon, Martin P Loeb, and Tashfeen Sohail. A framework for using insurance for cyber-risk management. *Communications of the ACM*, 46(3):81–85, 2003.
- [3] Jay Kesan, Ruperto Majuca, and William Yurcik. Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study. In *Workshop on the Economics of Information Security*, 2005.
- [4] Jean-Chrysostome Bolot and Marc Lelarge. A new perspective on internet security using insurance. In *INFOCOM Conference on Computer Communications*, pages 1948–1956. IEEE, 2008.
- [5] Daniel W Woods and Tyler Moore. Does insurance have a future in governing cybersecurity? *IEEE Security Privacy*, 18(1):21–27, Jan 2020.

- [6] Ulrik Franke. The cyber insurance market in Sweden. *Computers & Security*, 68:130–144, 2017.
- [7] Shauhin A Talesh. Data breach, privacy, and cyber insurance: How insurance companies act as “compliance managers” for businesses. *Law & Social Inquiry*, 43(2):417–440, 2018.
- [8] Jason R.C. Nurse, Louise Axon, Arnau Erola, Ioannis Agraftotis, 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 (Cyber SA)*. IEEE, 2019.
- [9] Ronald Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.
- [10] Oliver E Williamson. *The Mechanisms of Governance*. Oxford University Press, 1996.
- [11] Rainer Böhme and Galina Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Workshop on the Economics of Information Security*, 2010.
- [12] Kenneth J Arrow. Uncertainty and the welfare economics of medical care. *The American Economic Review*, 53(5):941–973, 1963.
- [13] Isaac Ehrlich and Gary S Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, 1972.
- [14] Oliver E Williamson. Transaction cost economics. *Handbook of Industrial Organization*, 1:135–182, 1989.
- [15] Herbert A Simon. *Models of Man; Social and Rational*. Wiley, 1957.
- [16] Herbert A Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129, 1956.
- [17] Karl N Llewellyn. What price contract—an essay in perspective. *Yale Law Journal*, 40:704, 1930.
- [18] Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. In Moore, Pym, and Ioannidis, editors, *Economics of Information Security and Privacy*.
- [19] Fabio Massacci, Joe Swierzbinski, and Julian Williams. Cyberinsurance and public policy: Self-protection and insurance with endogenous adversaries. In *Workshop on the Economics of Information Security*, 2017.
- [20] Mohammad Mahdi Khalili, Mingyan Liu, and Sasha Romanosky. Embracing and controlling risk dependency in cyber-insurance policy underwriting. *Journal of Cybersecurity*, 5(1):tyz010, 2019.

- [21] Mohammad Mahdi Khalili, Xueru Zhang, and Mingyan Liu. Effective premium discrimination for designing cyber insurance policies with rare losses. In *International Conference on Decision and Game Theory for Security*, volume 11836, pages 259–275. Lecture Notes in Computer Science, Springer, 2019.
- [22] Ganbayar Uuganbayar, Artsiom Yautsiukhin, Fabio Martinelli, and Fabio Massacci. Optimisation of cyber insurance coverage with selection of cost effective security controls. *Computers & Security*, 101:102121, 2021.
- [23] Sasha Romanosky, Andreas Kuehn, Lillian Ablon, and Therese Jones. Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity*, 5(1):tyz002, 2019.
- [24] Daniel W Woods, Ioannis Agraftotis, Jason RC Nurse, and Sadie Creese. Mapping the coverage of security controls in cyber insurance proposal forms. *Journal of Internet Services and Applications*, 8(1):8, 2017.
- [25] Daniel W. Woods and Rainer Böhme. SoK: Quantifying cyber risk. In *IEEE Symposium on Security and Privacy*, pages 909–926, Oakland, CA, May 2021.
- [26] Gregory Falco, Martin Eling, Danielle Jablanski, Matthias Weber, Virginia Miller, Lawrence A Gordon, Shaun Shuxun Wang, Joan Schmit, Russell Thomas, Mauro Elvedi, et al. Cyber risk research impeded by disciplinary barriers. *Science*, 366(6469):1066–1069, 2019.
- [27] Aron Laszka, Emmanouil Panaousis, and Jens Grossklags. Cyber-insurance as a signaling game: Self-reporting and external security audits. In *Proceedings of the 9th Conference on Decision and Game Theory for Security (GameSec 2018)*, volume 11199. Lecture Notes in Computer Science, Springer, 2018.
- [28] Sakshyam Panda, Daniel W Woods, Aron Laszka, Andrew Fielder, and Emmanouil Panaousis. Post-incident audits on cyber insurance discounts. *Computers & Security*, 87:101593, 2019.
- [29] Savino Dambra, Leyla Bilge, and Davide Balzarotti. SoK: Cyber insurance—technical challenges and a system security roadmap. In *IEEE Symposium on Security and Privacy*, pages 293–309, 2020.
- [30] American International Group (AIG). Claims Intelligence Series. [Online; accessed 20-Nov-2020].
- [31] Louise Axon, Arnau Erola, Ioannis Agraftotis, Michael Goldsmith, and Sadie Creese. Analysing cyber-insurance claims to design harm-propagation trees. In *2019 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (Cyber SA)*. IEEE, 2019.

- [32] Josephine Wolff and William Lehr. Roles for policy-makers in emerging cyber insurance industry partnerships. In *46th Research Conference on Communication, Information and Internet Policy (TPRC 46)*, 2018.
- [33] Nicholas Kaldor. A model of economic growth. *The Economic Journal*, 67(268):591–624, 1957.
- [34] Bryan A Garner. *Black’s law dictionary (Ninth Edition)*. West Group St. Paul, MN, 2009.
- [35] Advisen Ltd. Advisen’s Cyber Guide, 2019. [Online; accessed 20-Dec-2020].
- [36] Tony Vila, Rachel Greenstadt, and David Molnar. Why we can’t be bothered to read privacy policies models of privacy economics as a lemons market. In *5th International Conference on Electronic commerce*, pages 403–407. ACM, 2003.
- [37] Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799):610–613, 2006.
- [38] Cormac Herley and Dinei Florêncio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In Moore, Pym, and Ioannidis, editors, *Economics of Information Security and Privacy*, pages 33–53. Springer, 2010.
- [39] John Wadleigh, Jake Drew, and Tyler Moore. The e-commerce market for “lemons”: Identification and analysis of websites selling counterfeit goods. In *24th International Conference on World Wide Web*, pages 1188–1197, 2015.
- [40] Daniel W Woods and Tyler Moore. Cyber warranties: market fix or marketing trick? *Communications of the ACM*, 63(4):104–107, 2020.
- [41] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. Computer security incident handling guide. *NIST Special Publication*, 800(61):1–147, 2012.
- [42] Stan J Liebowitz and Stephen E Margolis. Path dependence, lock-in, and history. *Journal of Law, Economics, & Organization*, 11(1):205–226, 1995.
- [43] W Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394):116–131, 1989.
- [44] Keman Huang, Michael Siegel, and Stuart Madnick. Systematically understanding the cyber attack business: A survey. *ACM Computing Surveys (CSUR)*, 51(4):1–36, 2018.
- [45] Benjamin Collier, Richard Clayton, Alice Hutchings, and Daniel Thomas. Cybercrime is (often) boring: maintaining the infrastructure of cybercrime economies. In *Workshop on the Economics of Information Security*, 2020.

- [46] Rolf Van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganan, Bram Klievink, Nicolas Christin, and Michel Van Eeten. Plug and prey? Measuring the commoditization of cybercrime via online anonymous markets. In *27th USENIX Security Symposium*, pages 1009–1026, 2018.
- [47] Daniel W Woods and Andrew C Simpson. Policy measures and cyber insurance: a framework. *Journal of Cyber Policy*, 2(2):209–226, 2017.
- [48] European Union Agency for Network and Information Security (ENISA). Commonality of risk assessment language in cyber insurance.
- [49] US Department of Homeland Security. Insurance for Cyber-Related Critical Infrastructure Loss: Key Issues. 2014. [Online; accessed 4-March-2021].
- [50] US Department of Homeland Security. Cyber Risk Culture Roundtable Readout Report. 2013. [Online; accessed 4-March-2021].
- [51] US Department of Homeland Security. Healthcare and Cyber Risk Management: Cost/Benefit Approaches. 2014. [Online; accessed 4-March-2021].
- [52] UK Cabinet Office. UK Cyber Security: the Role of Insurance. [Online; accessed 7-June-2020].
- [53] Organisation for Economic Co-operation and Development. Supporting an Effective Cyber Insurance Market. 2017. [Online; accessed 7-June-2021].
- [54] Cyberspace Solarium Commission. The Cyberspace Solarium Commission Report. 2020. [Online; accessed 4-March-2021].
- [55] Daniel W Woods and Jessica Weinkle. Insurance definitions of cyber war. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 45:639–656, 2020.
- [56] Jan Martin Lemnitzer. Why cybersecurity insurance should be regulated and compulsory. *Journal of Cyber Policy*, in print, 2021.
- [57] Josephine Wolff. *Cyber-insurance Policy: Rethinking International Risk for the Internet Age*. Cambridge, MA: MIT Press, forthcoming 2022.
- [58] Shauhin A Talesh and Bryan Cunningham. The technologization of insurance: An empirical analysis of big data and artificial intelligence’s impact on cybersecurity and privacy. *Utah Law Review*, in print, 2021.
- [59] Benjamin Edelman. Adverse selection in online “trust” certifications and search results. *Electronic Commerce Research and Applications*, 10(1):17–25, 2011.
- [60] Jonathan Michael Spring. *Human decision-making in computer security incident response*. PhD thesis, University College London, 2019.

- [61] Howard Erlanger, Bryant Garth, Jane Larson, and Elizabeth Mertz. Is it time for a new legal realism. *Wisconsin Law Review*, pages 335–365, 2005.
- [62] W. S. Baer and A. Parkinson. Cyberinsurance in IT security management. *IEEE Security Privacy*, 5(3):50–56, 2007.
- [63] Angelica Marotta, Fabio Martinelli, Stefano Nanni, Albina Orlando, and Artsiom Yautsiukhin. Cyber-insurance survey. *Computer Science Review*, 14:35–61, 2017.
- [64] Adam Young and Moti Yung. Cryptovirology: Extortion-based security threats and countermeasures. In *IEEE Symposium on Security and Privacy*, pages 129–140, Oakland, CA, 1996.
- [65] Alexandre Gazet. Comparative analysis of various ransomware virii. *Journal in Computer Virology*, 6(1):77–90, 2010.
- [66] Amin Kharraz, William Robertson, Davide Balzarotti, Leyla Bilge, and Engin Kirda. Cutting the gordian knot: A look under the hood of ransomware attacks. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, volume 9148, pages 3–24. Lecture Notes in Computer Science, Springer, 2015.
- [67] Amin Kharaz, Sajjad Arshad, Collin Mulliner, William Robertson, and Engin Kirda. UNVEIL: A large-scale, automated approach to detecting ransomware. In *25th USENIX Security Symposium*, pages 757–772, 2016.
- [68] Nolen Scaife, Henry Carter, Patrick Traynor, and Kevin RB Butler. Cryptolock (and drop it): stopping ransomware attacks on user data. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 303–312. IEEE, 2016.
- [69] Eugene Kolodenker, William Koch, Gianluca Stringhini, and Manuel Egele. Paybreak: Defense against cryptographic ransomware. In *ACM Asia Conference on Computer and Communications Security*, pages 599–611, 2017.
- [70] Aron Laszka, Sadegh Farhang, and Jens Grossklags. On the economics of ransomware. In *8th Conference on Decision and Game Theory for Security*, volume 10575, pages 397–417. Lecture Notes in Computer Science, Springer, 2017.
- [71] Edward Cartwright, Julio Hernandez Castro, and Anna Cartwright. To pay or not: game theoretic models of ransomware. *Journal of Cybersecurity*, 5(1):tyz009, 2019.
- [72] Debabrata Dey and Atanu Lahiri. Should we outlaw ransomware payments? In *54th Hawaii International Conference on System Sciences*, pages 6609–6617, 2021.

- [73] Asaf Lubin. The insurability of cyber risk. *Available at SSRN 3452833*, 2019.
- [74] Dan Sabbagh. Insurers ‘funding organised crime’ by paying ransomware claims. *The Guardian*, 24 Jan 2021.
- [75] Tom Baker. On the genealogy of moral hazard. *Texas Law Review*, 75(2):237, 1996.
- [76] Viviana A Zelizer. *Pricing the Priceless Child: The Changing Social Value of Children*. Princeton University Press, 1994.
- [77] Viviana A Zelizer. Human values and the market: The case of life insurance and death in 19th-century America. *American Journal of Sociology*, 84(3):591–610, 1978.
- [78] Patrick T Brandt, Justin George, and Todd Sandler. Why concessions should not be made to terrorist kidnappers. *European Journal of Political Economy*, 44:41–52, 2016.
- [79] Alexander Fink and Mark Pingle. Kidnap insurance and its impact on kidnapping outcomes. *Public Choice*, 160(3-4):481–499, 2014.
- [80] Willis Towers Watson. Decode Cyber Podcast - Episode 2: Ransomware, 2021. [Online; accessed 7-June-2021].
- [81] Anja Shortland. *Kidnap: Inside the Ransom Business*. Oxford University Press, 2019.
- [82] Coveware. Ransomware Payments Up 33% As Maze and Sodinokibi Proliferate in Q1 2020, 2021. [Online; accessed 4-Mar-2021].
- [83] Coveware. Ransomware Payments Fall as Fewer Companies Pay Data Exfiltration Extortion Demands, 2021. [Online; accessed 4-Mar-2021].
- [84] Stefan Laube and Rainer Böhme. Strategic aspects of cyber risk information sharing. *ACM Computing Surveys (CSUR)*, 50(5):1–36, 2017.
- [85] Jennifer Anne Carlson. The economics of fire protection: from the great fire of London to rural/metro. *Economic Affairs*, 25(3):39–44, 2005.

A Recruitment Advert



Daniel Woods

Cyber Security PhD, University of Oxford

4mo • Edited •



Why do a handful of firms dominate cyber insurance carriers' lists of preferred providers?

We're launching a project exploring this ecosystem! Looking to speak to anyone involved in who gets selected as a post-breach provider. That means IR firms, breach coach/counsel, insurers, brokers, reinsurers etc.

Please get in touch if you want to participate or share this post if you're interested in seeing results further down the line.

[#cyber](#) [#dfir](#) [#cyberinsurance](#) [#cyberlaw](#)

Project funding from the EU:

<https://lnkd.in/gsBydAc>




Figure 11: Advert used for recruitment.

B Interview Guidelines

We prepared the following set of questions before the interviews began. Our semi-structured approach meant we deviated from the questions in order to explore answers in more depth. Also, we would not ask a question if the participant had already answered it in another question.

B.1 IR Provider

The following guidelines were used for forensics firms.

General

- Could you describe your professional background.
- What kind of services do you provide?
- How do you interact with insurers, brokers or breach coaches in your role?
- How many members in your team? Experience?

Search

- Who has influence in deciding which service provider is chosen?
- Rank the influence of insurers, brokers, breach coaches and the client in choosing the IR firms
- Do you see any dysfunctional aspects of the IR services ecosystem?
- What percentage of your relationships involved the insurer making first contact?
- Can you quote prices before understanding the incident? (e.g hourly rate or fixed price)
- Under what circumstances would you share quotes?

Negotiation

- Talk me through a typical or specific example of a negotiation with service selector
- How would you go about evaluating a service provider's quality?
- Do service providers negotiate with insurers/breach coaches/clients? Along which lines?
- What kind of agreements are there between insurer and service provider?
- Is this written down?
- Clearly negotiation is about price, but what is price being traded off against?

- Who decides what level of investigation takes place, how much time etc
- How does your pricing for service selector agreements compare to clients you find independently?
- How do the services in insurer agreements compare to clients you find independently?
- Are the insurers' prices negotiable?
- What percentage is the price for insurers compared to normal work?
- How often is it renegotiated?

Monitoring

- What happens to the forensic report?
- Who monitors service quality?
- What kind of disputes arise between insurer and service provider?
- How are they resolved?

High-level

- Do you anticipate any trends over the next few years?
- Could IR services be automated?
- What is the role of triage in IR response? Who decides how much resources get assigned to each incident?

B.2 Service selector

The following guidelines were used for insurers and breach coaches.

General

- Could you describe your professional background.
- What kind of services do you hire?
- How do you interact with forensics firms, brokers or insurers/breach coaches in your role?
- How many members in your team? Experience?

Search

- Who has influence in deciding which service provider is chosen?
- Rank the influence of insurers, brokers, breach coaches and the client in choosing the IR firms

- Do you see any dysfunctional aspects of the IR services ecosystem?
- In what percentage of your relationships did you make first contact?
- Do you negotiate prices before understanding the incident? (e.g hourly rate or fixed price)

Negotiation

- Talk me through a typical or specific example of a negotiation with service provider
- How would you go about evaluating a service provider's quality?
- Do service providers negotiate with insurers/breach coaches/clients? Along which lines?
- What kind of agreements are there between insurer and service provider?
- Is this written down?
- Clearly negotiation is about price, but what is price being traded off against?
- Who decides what level of investigation takes place, how much time etc
- Are your prices negotiable?
- What percentage discount do you negotiate?
- How often is it renegotiated?

Monitoring

- What happens to the forensic report?
- Who monitors service quality?
- What kind of disputes arise between insurer and service provider?
- How are they resolved?

High-level

- Do you anticipate any trends over the next few years?
- What is the role of triage in IR response? Who decides how much resources get assigned to each incident?

C Panel Analysis




Figure 12: Size of the publicly accessible vendor shortlists including the firms outside the Top 20..