8

CONCLUSION: IS CYBER RISK DIFFERENT?

Viewed alongside the emergence of other types of insurance, the development of the cyberinsurance market over the past three decades has been both disarmingly rapid and surprisingly slow. The rapidity has been demonstrated most vividly in the vast array of different policies and products that insurers have begun offering linked to cyber risks in the span of just a few decades. Unlike car, flood, or fire insurance, cyberinsurance does not cover a single, coherent type of threat, and unlike CGL or property and casualty insurance it does not cover a particular, coherent set of damages. Instead, cyber risk insurance, in its various forms from stand-alone policies to add-on products, tries to tackle a range of different threats, from cybercrime and data breaches to network outages, user errors, and online extortion—and across that wide range of threats it also aims to encompass an astonishing number of different types of damage, from first-party costs, such as lost business, breach notifications, and ransom payments, to third-party costs tied to lawsuits and liability.

Trying to describe the cyberinsurance industry makes clear the extent to which cyberinsurance is fundamentally not a single thing but rather a range of different products that deal with computer-, data-, and network-related risks that intersect with any number of different threats and types of losses. This would make the entire endeavor of studying cyberinsurance as a topic seem almost foolish were it not for the fact that insurers have increasingly tried to establish it as a single, coherent market with dedicated policies and coverage specifically for cyber risk. To this end, they have also excluded cyber-related losses from their other coverage, steering customers toward stand-alone cyber policies instead. Even in their internal structure and organization, many insurers have set up dedicated cyber risk groups to develop these policies, in many cases leaving the cyberinsurance team siloed apart from the groups working on modeling and pricing other, related risks in different departments.

The cyberinsurance market has grown slower than many carriers anticipated, even in the aftermath of a series of high-profile cybersecurity incidents and data protection regulations which insurers had predicted would significantly boost sales. For instance, a 2015 PricewaterhouseCoopers report titled "Insurance 2020 & beyond: Reaping the Dividends of Cyber Resilience" projected that the cyberinsurance industry would triple between 2015 and 2020, reaching annual premiums of roughly $7.5 billion by 2020.[1] Instead, in 2020, the NAIC estimated that the US market for cyberinsurance was still under $4 billion in premiums, and that the take-up rate for cyber policies remained "relatively low" at 33 percent.[2] A slew of ransomware attacks and other cybersecurity incidents beginning in 2019 also reduced the sizeable profit margins that carriers had previously enjoyed on cyber risk policies. In 2019, Aon estimated that the loss ratio for US cyberinsurance policies increased by 10 percent, to approximately 45 percent, compared to 35 percent in 2018.[3] That meant that in the span of one year, carriers went from paying out roughly 35 cents in claims for each dollar of premiums collected to paying out 45 cents per dollar of premium payments—a significant change, particularly given the reputation cyberinsurance had acquired by then for being "more profitable for insurers than other lines of insurance," as one 2019 *ProPublica* article put it, comparing the 35 percent loss ratio for cyberinsurance in 2018 to the 62 percent loss ratio for property and casualty insurance coverage.[4]

But these changes have not deterred insurers from developing and marketing new policies and new partnerships to address cyber risks. This drive to sell cyberinsurance may stem in part from carriers' desire to land customers while the market is still relatively new and businesses have not yet committed to a carrier, with carriers counting on their own ability to refine the risk models and pricing later, as they collect more data and learn more about the nature of cyber risk and the best methods for reducing exposure. But that assumption—that with time and data it will be possible to tame cyber risk using the same tools and techniques that have been applied with such success to so many other kinds of risk—relies on the idea that cyber risks are fundamentally no different from robberies or floods or car accidents or kidnappings in that they can be modeled and priced in their own comprehensive, stand-alone policies. This is not the case.

What differentiates cyber risk from other types of risk is not simply its scale, or how quickly it has evolved, or the complexity of computer

networks, or the presence of determined and intelligent adversaries, or the uncertainty about how to mitigate these risks most effectively—though all of those characteristics undoubtedly do add to the considerable challenges of trying to craft cyberinsurance coverage. What makes cyber risk different is that it is not a single type of risk, that it extends to and interconnects nearly every other type of risk—from crime to liability to property and casualty losses—in ways so unpredictable and unprecedented that it is hard to imagine these actuarial complexities being captured simply by the collection of more data or the use of more sophisticated modeling tools. These challenges of scale and interconnection echo, to some extent, the complexities that insurers have faced in covering growing environmental risks. Daunted by the potentially massive consequences of these risks for all forms of natural disaster and property coverage, insurers have at various times tried to limit their environmental liability by refusing coverage, raising rates for coverage, and even engaging with policymakers on initiatives such as signing an accord with the United Nations to address climate change. Not all of these approaches were necessarily constructive for actually preventing climate change. As Haufler points out, "the unavailability or high costs of insurance may simply mean that a lot of business will be uninsured; when accidents occur, someone else will have to pay for the cleanup, which often comes down to public money."[5]

At least for the time being, there appears to be no shortage of insurers willing to sell cyberinsurance policies at rapidly rising prices for businesses of all sizes. It's not necessarily clear, however, that those policies actually cover the range of risks that policyholders believe or expect them to. Those unmet expectations are partly a function of the lack of standardized policy templates or clarity around exceptions, but they are also tied to the fact that, unlike environmental risks that correspond to a fairly clear and well-understood set of natural disasters, neither carriers nor policyholders are necessarily able to anticipate the kinds of cyber risks that will emerge even one or two years into the future. From a policyholder's perspective, that would seem to make cyber risks a good candidate for a broad all-risk policy, in the style of property and casualty insurance, that promises to cover any type of risk other than those explicitly excluded by the carrier, relieving the policyholder of anticipating all the possible risks they may require coverage for in the future. But, for exactly the same reasons, insurers have been reluctant to offer overly broad all-risk cyber policies, and have

chosen instead to tailor narrow add-on products to existing coverage types and craft stand-alone named peril cyber policies that cover only a specific set of types of losses and liability. Over time, that set of losses has grown significantly larger as insurers have added more types of coverage to their stand-alone policies, particularly for first-party losses. That expansion of stand-alone cyber policies has been paralleled by a growing tide of resistance to policyholders claiming cyber-related losses under other types of coverage, even when that coverage includes riders and other add-on products specifically designed to cover computer-related losses.

The legal disputes over denied CGL, crime, and property and casualty insurance claims for cyber-related losses have helped clarify some of the ambiguities about what different types of insurance do and do not cover and have also, at times, reinforced the idea that there are many different, competing interpretations of what constitutes a cyber risk. For instance, while there is clear consensus that CGL policies do not apply to data breach litigation in most cases, there is much more uncertainty around whether incidents involving phishing emails are acts of computer fraud or not, or what constitutes a warlike act in cyberspace. In some ways, that uncertainty has been productive, driving carriers to clarify the language in their policies and exceptions. At the same time, however, it may also dissuade would-be customers from purchasing pricy coverage that they fear might not actually apply in the event of a significant cybersecurity incident. This, then, can lead to an outcome similar to the one Haufler observed for environmental risk insurance—companies choosing not to purchase coverage so that the costs of cyberattacks end up being borne by the public sector or individual victims.

Reliance on public funding to pay for cybersecurity incidents is not necessarily a terrible outcome—in fact, it is precisely what some insurers are lobbying for when they talk about extending TRIA to cyberattacks—but it does leave insurers in a less powerful position to enforce security standards and controls across a large customer base. For cyberinsurance to serve as an effective form of cybersecurity governance, insurers must be able to identify and incentivize policyholders to implement preventive measures that actually reduce the private sector's exposure to cyber risks, rather than just functioning as a form of compensation and risk pooling for victims of security incidents. It is striking that the insurance industry has, thus far, demonstrated so little progress on that front. Despite all the partnerships with security firms and the years of collected claims data, insurers seem to have

no greater insight into how to reduce a policyholder's risk exposure or prevent cybersecurity incidents than they did when the market emerged in the late 1990s. Policyholders are still vetted in largely cursory ways, according to brief questionnaires that typically yield little insight into an organization's technical defenses and have even less impact on their premiums.

Government interest in cyberinsurance has been predicated in large part on the notion that insurers will be able to reduce policyholders' exposure to cyber risk. As early as 2011, the United States Department of Commerce Internet Policy Task Force referred to cybersecurity insurance as a potentially "effective, market-driven way of increasing cybersecurity."[6] The following year, the DHS speculated it could "help reduce the number of successful cyber attacks by promoting widespread adoption of preventative measures, encouraging the implementation of best practices by basing premiums on an insured's level of self-protection, and limiting the level of losses that companies face following a cyber attack."[7] Nearly a decade later, the only one of those goals that insurers seem even close to being able to achieve is that last one: limiting third-party losses, post-breach, by providing policyholders with immediate incident response resources and legal counsel. But while reducing the amount of data breach–related litigation may significantly decrease the costs associated with those breaches for the companies in question and, by extension, their insurers, it's not clear that this actually increases cybersecurity for anyone, much less reduces the number of successful cyberattacks.

Reducing risk exposure is not the sole purpose of insurance. In some cases, it's not even the primary purpose, particularly when—as in the case of cyberinsurance—carriers find themselves unable to assess the risks faced by policyholders. Kenneth Abraham traces the development of workers' compensation programs designed to guarantee that employees would receive compensation for harm that befell them from accidents at work. The proposals for these programs focused more on ensuring compensation for victims than on accident prevention, Abraham argues. Carriers offering employers' liability policies in the early twentieth century struggled to even figure out what safeguards companies provided to limit accidents in the workplace and whether their customers complied with safety standards. "Travelers' own company history recounted the difficulties it encountered in getting policyholders to make safety changes," Abraham writes, citing an early Travelers liability insurance inspector who said of his experience in the company's

official history, "We enjoyed little cooperation and much downright antagonism. The boss had no interest in the elimination of the danger, and the workers themselves had become so used to conditions that they resisted change."[8] Similarly, New York's influential 1910 Wainwright Commission Report on workers' compensation made a "passing reference . . . to the potential of a workers' compensation system to reduce the incidence of accidents" but, Abraham points out, "the Report noted at the outset that the Commission had not yet been able to address the causes and prevention of accidents, promising to address these issues in a subsequent Report." He concludes: "A Report that recommends the enactment of workers' compensation before it has had the chance to address the causes and prevention of accidents must be understood to be concerned primarily with other issues."[9]

Much like the Wainwright Commission and the Travelers' insurance inspectors one hundred years before them, today's policymakers and carriers have not yet really been able to address the causes and prevention of cybersecurity incidents. For all the published frameworks, catalogs of security controls, and lists of best practices, there is no strong empirical evidence of what defenses are most effective at reducing cyber risk or even clear consensus on how to measure the impact of different security controls. Nothing has made that clearer than the unwillingness of insurers to make significant adjustments to premium prices based on their customers' security postures. And yet, unlike the Wainwright Commission, policymakers working on cyberinsurance have repeatedly lauded it as a means of helping prevent cybersecurity losses. Indeed, many government discussions appear to assume that the best way to address the lack of empirical evidence for the effectiveness of different cybersecurity measures is by building a robust cyberinsurance market that can collect and analyze the needed data. If that turns out not to be the case, then the cyberinsurance market may continue to function primarily as a means of sharing losses rather than preventing them, serving to pool premiums from a wide array of companies and using that money to compensate the victims of breaches, outages, and other computer compromises. That, in itself, could be a worthwhile goal, but the concern is that cyberinsurance, if it doesn't succeed in bolstering security standards, could actually lead to the deterioration of policyholders' security practices due to moral hazard. Even worse, if cyberinsurance means that extortion payments become a widely accepted and routinized part of doing business, then this type of coverage will contribute to the growth of the cybercrime market by underwriting

extortion payments that both indirectly encourage and directly fund further criminal activity.

Nearly every challenge that insurers currently face in trying to model and price cyberinsurance reflects a problem they have encountered—and in many cases, solved—before, in the history of insurance. Selling car insurance required carriers to collect data about the evolving risks of a new and changing technology. To offer crime insurance, insurers had to take into consideration the actions of an intelligent adversary who can adapt to preventive countermeasures. Developing kidnapping and ransom policies meant dealing with the potential unintended consequences of making direct payments to criminals and thereby encouraging copycats. Designing terrorism coverage forced insurers to face the possibility of catastrophic, accumulated risk. What those types of insurance have in common—and do not share with cyberinsurance—is that they cover a coherent and relatively stable set of risks.

Car accidents, crimes, ransom, terrorism—none of those risks has changed dramatically in nature in the past several decades except for their computer-based components. The task that falls to insurers in developing cyberinsurance, then, is not just to model and understand a new class of risk but also to remodel and rethink nearly every other existing class of risk they cover. No wonder they have gone to such lengths to try to exclude many cyber-related claims from their customers' existing insurance and tried to shift as much cyber-related risk coverage as possible into isolated stand-alone cyber policies. That is the approach the insurance industry has taken with nearly every new set of risks it has expanded to cover. It allows carriers to continue to rely on their core business and products while exploring a new area, but at the same time it leaves them further entrenched in the idea that each of these classes of risk is distinct and distinguishable.

Looking ahead, cyber risks will only become increasingly intertwined with the existing classes of risks insurers cover. Autonomous vehicles will require carriers to rethink auto insurance, buildings furnished with Internet-connected heating and cooling systems, fire sprinklers, and security cameras will change property insurance. Devices that can constantly monitor users' heart rates, activity levels, and other health indicators may similarly transform the field of health insurance. In some cases, these new technologies may enable insurers to monitor their policyholders more closely and require or recommend more stringent, high-tech safeguards against risks like car accidents, robberies, or heart attacks. But, inevitably, even as technologies

like self-driving cars, security cameras with facial recognition capabilities, or health trackers may help reduce our exposure to some of these threats, they will also create new risks and introduce new avenues of attack via the complicated systems they connect to our cars, homes, and bodies.

Designing car insurance for autonomous vehicles won't just require adjusting the existing models and policies, it will require radically reimagining them for a set of risks we know very little about, such as computer vision errors and vulnerabilities in car software systems. Beyond just trying to collect enough data to understand how frequently these types of risks occur and what their financial impacts are, insurers and policymakers will also have to rethink questions related to liability: who is responsible for car accidents that occur because of malicious software compromises or faulty machine learning algorithms? The introduction of computers and computer networks to existing systems doesn't just create new risks for those systems, it also introduces a new set of stakeholders and intermediaries who are involved in designing the relevant software and hardware, connecting those legacy systems to a larger network of computers, and then monitoring those connections to restrict malicious activity. All of these stakeholders, in addition to those who were already involved—the car manufacturer and the drivers, for instance—play a role in mitigating risks that are in some way connected to computers and are therefore important for thinking about effective and comprehensive liability regimes.

Insurers will probably look to the courts, and perhaps also to regulators, to help decide how these complicated liability issues will be resolved. This has been true in the past, as insurers have taken their cues about what types of liability coverage to offer and to whom from civil lawsuits and the resulting rulings. Reflecting on the history of liability insurance, Abraham argues that "tort law continually seeks an available source of recovery, creating or expanding the liability of individuals and businesses that are likely to be covered by or have access to liability insurance. And liability insurance has usually responded, by creating new forms of insurance to meet the new liabilities when such insurance was not already available."[10] But for there to be civil lawsuits about who is liable for autonomous vehicle accidents there first have to be enough such accidents for someone to sue, and it's not clear that people will begin driving—or even selling—autonomous vehicles in any significant numbers until there is adequate insurance in place to protect them from liability. In other words, the typical cycle of insurers waiting

for courts to dictate new liability regimes and then crafting policies to fit those regimes may not work for certain types of cyber risks associated with activities like driving where insurance is expected, if not required. If insurers are unable to get a handle on coverage for cyber risks of all varieties, that could significantly slow, or even prevent, the process of people and business beginning to adopt new technologies available to them.

Another concern is the possibility that emerging cyber risks will lead to a narrowing of insurance coverage rather than an expansion. Already, cyber-related losses are being explicitly excluded from many types of insurance but, for the most part, those exclusions are balanced by the development of new cyber risk policies that cover much of what is excluded from carriers' other coverage. However, as they encounter new types of risk, insurers may decide there are some kinds of cyber risks they simply do not see themselves being able to cover. Abraham points out that while insurers often respond to court rulings that create or expand liability in new areas by expanding their coverage offerings, this is not always the case. "Sometimes insurers cannot, or will not, provide insurance against a new liability," he writes.

As an example, Abraham points to the expansion of pollution cleanup liability in the 1980s, following the passage of the Comprehensive Environmental Response, Compensation and Liability Act (CERCLA). In this case, instead of leading to broader insurance coverage for pollution cleanup, the new, stricter liability regime "led to the virtual disappearance of pollution liability insurance rather than to its expansion. Expansive judicial interpretations of insurance policies that had seemed to insurers to provide only limited pollution liability insurance to their policyholders eventually caused the insurance industry to insert an 'absolute' pollution exclusion into subsequently issued policies."[11] It is not hard to imagine similar exclusions for certain types of cyber risks emerging in the wake of expansive judicial interpretations of insurance policies that insurers thought offered only limited cyber coverage. For instance, if the courts rule that the property policies held by Mondelez and Merck actually do cover the damages caused by NotPetya, the CERCLA example suggests that insurers might decide to reduce the scope of their cyber coverage for certain types of risks rather than expanding it.

If insurers do continue to expand their coverage of cyber risks, there is no shortage of looming threats and problems on the horizon from which businesses—and perhaps even, eventually, individuals—will be eager to protect themselves. Insurance has only barely begun to grapple with the

risks presented by the Internet of Things and the proliferation of artificial intelligence (AI) and the use of machine learning algorithms for decision making. The risks associated with Internet of Things devices are likely to be entangled with existing insurance products, including auto insurance and property insurance, while the risks associated with AI may present more opportunities for entirely new forms of coverage. In an article titled "The Case for AI Insurance," Ram Shankar Siva Kumar and Frank Nagle point out that "AI failures resulting in business interruption and breach of private information are most likely covered by existing cyber insurance, but AI failures resulting in brand damage, bodily harm, and property damage will not likely be covered by existing cyber insurance."

Kumar and Nagle propose that companies should be taking stock of the safety and security of their AI systems and talking to their insurers about potential coverage for both intentional and unintentional failures of those systems. "We believe that AI insurance will first be available via major insurance carriers as bespoke insurers may not have sufficient safety nets to invest in new areas," they predict. "From a pricing perspective, using the past cyber insurance market as a template, businesses can expect stringent requirements when AI insurance is introduced to limit the insurance provider's liability with rates cooling off as the AI insurance market matures."[12] In fact, the short history of the past cyberinsurance market suggests a more complicated trajectory than just falling rates and less stringent liability limitations over time. Historical parallels might predict a gradual shuffling of different cyber risks from add-on products to stand-alone policies, accompanied by an expanding set of exclusions and no clear decrease in premium payments.

It's not surprising that insurers would look to excise cyber risks from non-cyber-specific policies and isolate them in stand-alone cyberinsurance policies in order to protect their existing core products from the uncertainty and unpredictability of cyber risk. But that isolation can also be counterproductive, for both carriers and their customers, when it gives credence to the idea that computer networks and data pose a distinct, definable set of risks that can be separated from the other categories of risk that insurers cover and policyholders face. Some cyber risks, like data breaches, AI algorithm errors, and online extortion, may in fact be so new and so unrelated to other, existing coverage that it makes sense for them to be covered in stand-alone policies, but as computer networks are increasingly embedded in existing physical infrastructure and systems, many—perhaps most—of

the risks they present will belong under the same policies that already protect those domains.

This is what is most fundamentally new and different about cyber risk as compared to other types of risks that insurers have addressed in the past—not just that it can, at times, be more unpredictable or more catastrophic or more difficult to mitigate, but that it requires remodeling so many other categories of risks, in addition to creating a new class of insurance products for risks to entirely new kinds of infrastructure and operations. Insurers look to data collection to help shape their policies, but this is not a challenge that will diminish with time, as more data is collected and analyzed. Rather, it is a challenge that will only grow as computing technology continues to extend into new areas and applications. Moreover, part of the challenge of rethinking existing risk categories will involve acknowledging the increasing interconnectedness among them and the potential for a single attack to have significant impacts related to property damage, car accidents, liability, business interruption, data breaches, crime, and terrorism, simultaneously. In this regard, cyber risks may, in fact, render existing insurance risk categories more unpredictable, more catastrophic, and more difficult to mitigate than ever before.

Cybersecurity, like climate change, will require the involvement of regulators and policymakers to make an insurance market viable in the long term, and that involvement will probably not be limited to just serving as a data aggregator or financial backstop for the insurers. It may well require regulators to take an active role in requiring certain cybersecurity standards and controls—as the EU has already begun to do for critical infrastructure operators through the NIS Directive—rather than waiting for insurers to identify those safeguards themselves and screen policyholders for them. It may also require regulators to take a hardline stance on the coverage and payment of online extortion demands which benefit certain stakeholders, including both carriers and cyberinsurance policyholders, in the short term but inflict significant harm, long term, by funding criminal enterprises and driving increased cybercrime.

The idea that cybersecurity can be handled solely, or even primarily, through a market-driven approach led by insurers is fundamentally flawed—something that insurers themselves, to their credit, have been pointing out to policymakers for years. Policymakers, too, have shown greater willingness to regulate data protection, particularly where it involves individuals' personal information, and especially outside the United States, where cyberinsurance

remains relatively uncommon. Some elements of those regulations, particularly incident reporting requirements and cybersecurity certifications, seem aimed, at least in part, at helping insurers develop better cyberinsurance policies. Other components, such as data localization measures, may instead serve to enervate the global cyberinsurance industry. Whether or not the wave of data protection regulations around the world in the late 2010s will actually drive greater adoption of cyberinsurance in those countries remains to be seen, but at the very least such regulations suggest that a growing number of governments are abandoning the notion that cybersecurity is something that can be solved by the private sector alone.

That should not diminish what the cyberinsurance industry has accomplished in developing a wide array of offerings for first- and third-party coverage for cyber risks all in the span of less than three decades, however. In the early 2000s, carriers significantly expanded available coverage for cyber-related losses to include insurance for network outages, restoration of encrypted systems, cryptocurrency-based crimes, and social engineering. That progress comes despite the ambiguity of some of those policies and the contentious legal disputes over what they do and don't apply to, and despite insurers' apparent inability to identify effective security controls and unwillingness to share claims data with their competitors. This expansion has been driven by demand from policyholders, but insurers have met this demand at considerable long-term financial risk to themselves since very little is known about how these threats will evolve over time or how courts will interpret the coverage and exclusions in these policies in light of future incidents.

As insurers continue to expand their cybersecurity coverage, they should also consider expanding the boundaries of how they define and conceptualize cyber risk within their organizational structures and underwriting categories. This means acknowledging the complicated and extensive connections between cyber risk and other coverage areas and crafting policies that recognize and reflect those connections. In the past, when a significant new type of risk has emerged, whether in the form of a novel type of legal liability or an innovative technology, the insurance sector has developed new products to cover those risks. When it comes to tackling cyber risk, however, the most important thing insurers can do is reinvent their old policies, rather than write new ones. Not all risks are cyber risks, but, increasingly, all types of risk have cyber components that insurers and their policyholders ignore or isolate at their peril.

Cyberinsurance Policy

Rethinking Risk in an Age of Ransomware, Computer Fraud, Data Breaches, and Cyberattacks

By: Josephine Wolff

Citation:

Cyberinsurance Policy: Rethinking Risk in an Age of Ransomware, Computer Fraud, Data Breaches, and Cyberattacks
By: Josephine Wolff
DOI: 10.7551/mitpress/13665.001.0001
ISBN (electronic): 9780262370752
Publisher: The MIT Press
Published: 2022

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

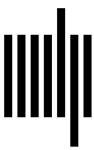

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Bembo by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Wolff, Josephine, author.
Title: Cyberinsurance policy : rethinking risk in an age of ransomware, computer fraud, data breaches, and cyberattacks / Josephine Wolff.
Description: Cambridge, Massachusetts : The MIT Press, [2022] | Series: Information policy series | Includes bibliographical references and index.
Identifiers: LCCN 2021045988 | ISBN 9780262544184 (paperback)
Subjects: LCSH: Computer insurance. | Computer security—Management. | Cyberspace—Security measures—Management. | Computer crimes—Prevention. | Risk management.
Classification: LCC HG9963.5 .W65 2022 | DDC 658.4/78—dc23/eng/20220114
LC record available at https://lccn.loc.gov/2021045988