Will Catastrophic Cyber-Risk Aggregation Thrive in the IoT Age? A Cautionary Economics Tale for (Re-)Insurers and Likes

RANJAN PAL and ZIYUAN HUANG, University of Michigan, Ann Arbor, USA
SERGEY LOTOTSKY, University of Southern California, USA
XINLONG YIN and MINGYAN LIU, University of Michigan, Ann Arbor, USA
JON CROWCROFT, University of Cambridge, UK, Alan Turing Institute, UK
NISHANTH SASTRY, University of Surrey, UK, King's College London, UK
SWADES DE, Indian Institute of Technology Delhi, India
BODHIBRATA NAG, Indian Institute of Management Calcutta, India

Service liability interconnections among networked IT and IoT-driven service organizations create potential channels for cascading service disruptions due to modern cybercrimes such as DDoS, APT, and ransomware attacks. These attacks are known to inflict cascading catastrophic service disruptions worth billions of dollars across organizations and critical infrastructure around the globe. Cyber-insurance is a risk management mechanism that is gaining increasing industry popularity to cover client (organization) risks after a cyber-attack. However, there is a certain likelihood that the nature of a successful attack is of such magnitude that an organizational client's insurance provider is not able to cover the multi-party aggregate losses incurred upon itself by its clients and their descendants in the supply chain, thereby needing to re-insure itself via other cyber-insurance firms. To this end, one question worth investigating in the first place is *whether an ecosystem comprising a set of profit-minded cyber-insurance companies, each capable of providing re-insurance services for a service-networked IT environment, is economically feasible to cover the aggregate cyber-losses arising due to a cyber-attack.* Our study focuses on an empirically interesting case of *extreme heavy tailed cyber-risk distributions* that might be presenting themselves to cyber-insurance firms in the modern Internet age in the form of catastrophic service disruptions, and could be a possible standard risk distribution to deal with in the near IoT age. Surprisingly, as a negative result for society in the event of such catastrophes, *we prove via a game-theoretic analysis* that it *may not be economically incentive compatible*, **even under i.i.d. statistical**

This work has been supported by the NSF under grants CNS-1616575, CNS-1939006, and ARO W911NF1810208. The authors acknowledge small text usage from multiple sources according to a "fair-use" fashion. The authors further acknowledge Ibragimov 2009 for re-using some of their (a) notations, and (b) proof constructs of their results in Section 6 for the purposes of consistency and self-sufficiency towards building an amalgamated analysis framework for IoT cyber-risk scenarios.
Authors' addresses: R. Pal, Z. Huang, X. Yin, and M. Liu EECS, University of Michigan Ann Arbor, 1301 Beal Avenue, Ann Arbor, Michigan, 48109, USA; emails: {palr, ziyuanh, connory, mingyan}@umich.edu; S. Lototsky, Department of Mathematics, University of Southern California, Kaprielian Avenue, Los Angeles, California, 90089, USA; email: lototsky@usc.edu; J. Crowcroft, Computer Laboratory, University of Cambridge, CB3 0FD, Cambridge, UK; email: jon.crowcroft@cl.cam.ac.uk; N. Sastry, Department of Computer Science, University of Surrey, Guildford, Surrey, GU2 7XH; email: n.sastry@surrey.ac.uk; S. De, Department of Electrical Engineering, Indian Institute of Technology Delhi, Hauz Khas, New Delhi, 110016, India; email: swadesd@ee.iitd.ac.in; B. Nag, Operations Research and Management, Indian Institute of Management Calcutta, Kolkata 700104, India; email: bnag@iimcal.ac.in.

2158-656X/2021/05-ART17 $15.00
https://doi.org/10.1145/3446635

conditions on **catastrophic** cyber-risk distributions, for limited liability-taking risk-averse cyber-insurance companies to offer cyber re-insurance solutions *despite the existence of large enough market capacity to achieve full cyber-risk sharing*. However, our analysis *theoretically endorses* the popular opinion that **spreading i.i.d. cyber-risks** that are **not catastrophic** is an effective practice for aggregate cyber-risk managers, a result established theoretically and empirically in the past. A failure to achieve a working re-insurance market in critically demanding situations after catastrophic cyber-risk events strongly calls for centralized government regulatory action/intervention to promote risk sharing through re-insurance activities for the benefit of service-networked societies in the IoT age.

CCS Concepts: • **Security and privacy → Distributed systems security**; • **Mathematics of computing → Probability and statistics**;

Additional Key Words and Phrases: Cyber-risk, cyber-risk aggregation, re-insurance, extreme heavy-tail

ACM Reference format:
Ranjan Pal, Ziyuan Huang, Sergey Lototsky, Xinlong Yin, Mingyan Liu, Jon Crowcroft, Nishanth Sastry, Swades De, and Bodhibrata Nag. 2021. *Will Catastrophic Cyber-Risk Aggregation Thrive in the IoT Age?* A Cautionary Economics Tale for (Re-)Insurers and Likes. *ACM Trans. Manage. Inf. Syst.* 12, 2, Article 17 (May 2021), 36 pages.
https://doi.org/10.1145/3446635

1 INTRODUCTION

Global commerce is undergoing a profound digital transformation. As it becomes increasingly electronic and IoT driven (courtesy of the upcoming 5G technology), critical exposures in this sector are getting highly data driven. As a result, the majority of modern business and economic risks are subsequently becoming cyber in nature. More importantly, such cyber-risks are often networked and accumulate in a variety of different ways, thereby affecting many lines of business. As an example, commercial companies in diverse sectors such as automobiles, electronics, energy, finance, aerospace, etc., and their mutual trading relationships are characterized by systemic network linkages through major software providers (e.g., Oracle for DBMS support). A cyber-attack (e.g., an APT-driven zero day attack) motivated by a vulnerability in a software version can have a catastrophic cascading service disruption effect that might amount to net commercial losses worth billions of dollars across the various service sectors—a cyber-analog to the current human COVID19's[1] catastrophic impact on world business. As well-documented commercial cyber-attack examples in reality, the very recent *Mirai* DDoS (2016), *NotPetya* ransomware (2017), and *WannaCry* ransomware (2017) attacks due to compromise of large-scale IoT devices caused havoc among firms in various industries (having trading relationships among them) across the globe, resulting in huge financial losses for the firms due to them being deemed dysfunctional in providing service. The reader is referred to Coburn et al. [2018] and the appendix for additional examples of cyber-attacks capable of causing catastrophic systemic loss impacts.

Technically, the emergence of such large-impact cyber-attacks on the IoT terrain is not surprising to say the least. Based on Advisen's 2017 Annual Report, security is often not being built into the design of these IoT products with the rush to get them to market. Symantec's research on IoT security has shown that the state of IoT security is poor: (1) roughly 19% of all tested mobile apps used to control IoT devices did not use **Secure Socket Layer (SSL)** connections to the cloud; (2) approximately 40% of tested devices allowed unauthorized access to back-end systems;

[1]The recent COVID-19 infection outbreak can be thought of as an analog of a viral APT on humans. In practice, a cyber-COVID can be far more disastrous than the human COVID-19 as the former could stall (pervasive/ubiquitous) computing life, which the latter does not.

(3) around 50% did not provide encrypted firmware updates, if updates were provided at all; and (4) IoT devices usually had weak password hygiene, including factory default passwords; for example, adversaries use default credentials for the Raspberry Pi devices to compromise devices. The Dyn attack compromised less than 1% of IoT devices. By some accounts, millions of vulnerable IoT devices were used in a market with approximately 10 billion devices. *XiongMai Technologies*, the Chinese electronics firm behind many of the webcams compromised in the attack, has issued a recall for many of its devices. Based on this evidence, Shankar Somasundaram, senior director, Internet of Things, at *Symantec*, expects more of Dyn-like outages in the near future.

In the wake of major targeted corporate and critical infrastructure cyber-attacks (e.g., attacks on Sony, Target, the Ukraine power grid) in the past half decade, risk mitigation has become a top board-level concern across many organizations worldwide. As a result, transfer-based risk management products like cyber-insurance, which currently has a rapidly growing market[2] (*Source:* Betterley Annual Report, 2015 [Betterley 2015], Advisen annual report 2016), is a major go-to solution for the current corporate sector worldwide, in the event of a cyber-attack. Renowned cyber-security expert Bruce Schneier has popularly envisioned the current digital age to be one of incident response through third-party risk management solutions [Dambra et al. 2020]. Moreover, such solutions bear the socially beneficial promise (over their purely technical counterparts) of improving cyber-security by incentivizing users and organizations to voluntarily invest in appropriate amounts of the various elements of the portfolio [Lelarge and Bolot 2009a; Johnson et al. 2011; Anderson and Moore 2009].

1.1 Research Motivation

Market surveys suggest that demand for cyber-insurance significantly exceeds the capacity currently provided by the insurance industry. One primary reason that most insurers give for being cautious about expanding capacity is the *accumulation risk posed by cyber-threats.* A major fear among insurers here is that cyber-threats are inherently scalable and systemic through their spread via network inter-connectivity—a single malicious email generated by a botnet activity can result in an entire organization becoming dysfunctional with respect to the service it provides, and in turn potentially affect business services of all other organizations that depend on it.

In the event of cascading service disruptions due to a major cyber-attack, if all these organizations were to hold responsible their parent organization(s) on which they depend for providing services, it is quite likely that the insurance company of a certain root organization could need to bear the responsibility of covering a huge aggregate/accumulated risk of all or multiple organizations in the service chain via the design of re-insurance contracts [Millaire 2016]. Shouldering this responsibility clearly may not be aligned with satisfying the budget constraints and profit requirements of most risk-averse cyber-insurers open to the idea of providing re-insurance services, let alone cyber-forensic and risk data availability challenges they might need to overcome to implement accumulative coverage policies [Millaire 2016].

In practice, the idea of spreading aggregate cyber-risk among re-insurers is gaining traction [Coburn et al. 2018; Kessler 2014; LI 2018] for smart society settings whereby insurers covering aggregate cyber-risk of organizations in a given sector (e.g., manufacturing) wish to spread that risk among insurers of firms that are higher up in the supply chain (e.g., energy companies). *However, there is no mathematical analysis on the effectiveness of this idea for general individual cyber-risk distributions, without which aggregate cyber-risk managers may not have the confidence to scale their service markets.* Moreover, there may be significant differences in the cyber and non-cyber

[2]Recently, *Wired* magazine (2019) projected a huge USD 5.5 billion cyber-insurance market by 2030, with firms like HSB, part of Munich Re, and Bajaj Allianz starting to sell personal cyber-insurance.

re-insurance settings. More specifically, the benefit of unlikely large-impact systemic outcomes in the latter (as qualitatively stated in Kessler [2014]) may not apply to cyber-settings that involve a pervasive digital world with billions of IoT devices in operation. One could argue in light of recent empirical works [Biener et al. 2015; Eling and Wirfs 2019; Romanovsky 2013] that cyber-losses are less heavy tailed (longer tails an indicator of the likelihood of large-impact systemic outcomes) than generally perceived to be. *However, such results are clearly not definitive* given (1) a popularly known fact that there is a severe lack of reported rich cyber-loss data with detailed security and exposure indicator features to be able to generalize these empirical results in a principled fashion [Dambra et al. 2020], (2) contradictory claims by other works [McNeil et al. 2015], and (3) the un-precedented scale of (in-secure) pervasive IoT networks of the future (not accounted by these studies) that might easily contradict claims made in Biener et al. [2015], Eling and Wirfs [2019], and Romanovsky [2013]. *Consequently, we strongly vouch for considerable possibilities in the IoT age of potentially several hundred systemically important vendors that could be susceptible to concurrent and substantial business interruption in time and space.* This includes (among many others) at least eight DNS providers that service over 50,000 websites, and some of these vendors may not have the kind of security that exists within rich cloud providers like AWS that boast of multiple layers of security and data redundancy.

Our Focus - Our focus in this article is to conservatively analyze whether providing risk re-insurance services, in events of systemic/cascading cyber-loss impacts of considerable degree, is economically incentive compatible for a set of cyber-insurers that are open to the idea of providing cyber re-insurance services to their clients, i.e., other cyber-insurance firms. Our investigation is of immense importance to societal benefit in the age where IT-networked services are ubiquitous in nature, and complex cyber-attacks, both on commercial and critical infrastructures, are on the rise. A failure to properly manage risk via re-insurance services and/or a failure of re-insurance markets is really bad for society in general as cyber-insurance companies could individually opt out of covering aggregated cyber-risk, and a spiral effect could lead to the breakdown of a significant portion (if not entirely) of the cyber-insurance industry in the long run. This could (1) negatively impact overall cyber-security[3] in the network of organizations tied to one another via commercial liability relationships and (2) affect societal welfare whereby loss-induced depression could make organizations look forward to insurance firms for support but not find it when they need it most.

1.2 Research Contributions

We make the following research contributions in this article:

- —En route to proposing our system model, we first provide a mathematical intuition on why the nature of cyber-risk distributions is a prime determinant to the success of re-insurance markets for the aggregate cyber-risk-managing cyber-insurance industry (see Section 2). We rely on the empirical evidence of certain cyber-risk distributions being heavy tailed and tail dependent [Biener et al. 2015; Eling and Schnell 2020; Xu et al. 2018; Maillart and Sornette 2010] and show why such risk types make profit-minded risk-averse cyber-insurers go low on confidence to expand coverage markets, where coverage is on an aggregate sum of such heavy-tailed cyber-risks (see Section 2). We then propose our system model (see Section 3).
- —The model leads us to the design of a formal dynamic (single-round, two-stage) game-theoretic framework of re-insurance markets. We observe on analysis that the

[3]The reader is referred to Section 7 of the article, where we mention a list of papers that endorse the opinion that cyber-insurance can help improve network security.

re-insurance market Nash equilibrium (RMNE) is not unique and can result, even under insurer-ideal settings (i.i.d. heavy-tailed cyber-risks without tail-dependence),[4] in a socially disadvantageous stable state, i.e., equilibrium, where individual re-insurers do not find it incentive compatible to provide re-insurance or primary cyber-insurance services despite the pool of cyber re-insurers having enough spare capacity at the RMNE to cover aggregate cyber-losses (see Sections 3 and 4). *The proof that spreading* catastrophic*heavy-tailed cyber-risks (curtailed or otherwise) that are identical and independently distributed (i.i.d.), i.e., not tail dependent,* is not*an effective practice for aggregate cyber-risk managers is a surprising new facet that we unravel in this article via theory and contrasting to many existing results.* However, our analysis theoretically endorses the popular opinion that spreading i.i.d. cyber-risks that are *not catastrophic* is an effective practice for aggregate cyber-risk managers, a result established theoretically and empirically in the past by other works.

—We mathematically characterize the conditions under which the above-mentioned socially disadvantageous equilibrium outcomes can occur. As a main result, we observe that under certain heavy-tailed cyber-risks and in situations when individual cyber-insurers have limits on their coverage liability, the RMNE leads to non-incentive-compatible coverage outcomes where cyber-insurers are not inclined to provide primary cyber-insurance services, let alone re-insurance services. We validate our proposed theory behind such outcomes with an experimental evaluation conducted on network graphs resembling a supply chain network. Consequently, we provide some regulatory viewpoints about resolving, via the intervention of non-private government agencies, the aggregate risk coverage challenges arising from such socially disadvantageous equilibrium outcomes (see Sections 4, 5, 6, and 8).

The contributions complement our two recent efforts [Pal et al. 2020c, 2020a], where we conduct an orthogonal formal analysis on the *statistical incentive compatibility of providing cyber-reinsurance services,* the first of its kind, in IoT societies under complete general cyber-risk distribution types be it i.i.d./non-i.i.d. heavy tailed or otherwise. The analysis in these papers, unlike here, is oblivious to re-insurer rationality while making service feasibility decisions on risk spreading. However, they account for the *complete general nature* of statistical distributions in the decision-making process.

1.3 Relevance of Our Research to Management Information Systems (MIS)

Here, we draw the specific relevance of our research contribution to the domain of modern management information systems.

The Salient MIS Features - Organizational **information systems (ISs)** are interrelated components, i.e., technologies, people, and processes, working together to collect, process, store, and disseminate information to support decision making, coordination, control, analysis, and visualization in an organization. To this end, in our work, MIS is the means by which (1) information in networked (corporate) organizations is transmitted over the Internet/Intranet, (2) the software that displays the information such as Microsoft Excel, or (3) the data management systems that manage organizational data such as Oracle DB. Following the philosophy laid down in Nicholas Carr's Harvard Business Review article [Carr 2003], we assume the organizational MISs in the IoT age would treat information technology as a commodity that needs to be **effectively managed** to (1) **reduce organizational costs**, (2) ensure that organizations are **always running**, and (3) be as **cyber-risk free** as possible. *However, the management goal is multi-fold challenging today, simply*

[4]It is commonly known that heavy-tailed cyber-risks with tail dependence are troublesome to the insurance industry when they aggregate [Eling and Schnell 2020], whereas i.i.d. cyber-risks are suitable for loss coverage.

because IoT devices are often cheap and not very reliable to keep running always, and the security practices for such devices are too naive to be risk-free.

Increased Proliferation of IoT-Driven Businesses to Boost Profit - Despite high cybersecurity risks to run businesses with an increased reliance on IoT, the latter technology is all geared up to revolutionalize the former. To drive home this point, we take the retail sector as a motivating example business application for the IoT age, where it is forecasted that 80% of global retailers will deploy IoT solutions in the retail market by 2021. IoT applications are what allows retailers to raise productivity, improve customer experience, reduce costs, and increase sales. *Frontier Economics* estimates that increasing **machine-to-machine (M2M)** connections by 10% would generate a USD 2.26 trillion increase in the US GDP alone. Forward-thinking retailers are already reaping the benefits of deploying IoT use cases in retail. According to an Oracle report on the impact of IoT on customer excellence [Oracle 2019], 66% of surveyed retail executives state that IoT has already positively impacted their customer experience processes. Similarly 88% of the study's respondents state that using IoT in retail will provide better customer insight than any other data-gathering method. The research by Oracle cited above found that 50% of current IoT users report better insight into their customers' needs and preferences. It also found that 47% were now able to provide a better and more differentiated customer experience. Most retail stores have realized the potential of IoT to vastly improve their supply chain management, which is likely why global retail tech spending will grow 3.6% to reach almost USD 203.6 billion in 2019. Inventory management remains a headache for retailers. Inaccurate inventory tracking can cause overstocking, stockouts, and shrinkage, leading to retailers estimating that their current inventories are only 66% accurate. But IoT in retail can automate inventory visibility, thereby solving these problems for good. Checkout remains one of the most labor-intensive retail operations and an unpleasant process for customers. If a store is overcrowded, shoppers often decide to just leave. IoT solutions offer the opportunity to automate and personalize checkout. The improved checkout system can read tags on each item as customers leave and automatically charge the customer's mobile payment app. This personalization can lead to increasing revenue for the store. According to McKinsey's research [Bughin et al. 2015], 83% of customers say they would prefer if their shopping experience was personalized, while store revenues are reported to increase by 20% to 30% through effective personalization.

Managing Vulnerabilities in MIS - Increased reliance of businesses on IoT in the future implies more value to cyber-hackers after attack on ISs. From a security perspective, as already mentioned earlier, the cyber-threats to modern information systems include equipment failure, environmental disruptions, human or machine errors, and purposeful attacks that are often sophisticated, disciplined, well organized, and well funded. When successful, attacks on information systems can result in serious or catastrophic damage to organizational operations and assets, individuals, other organizations, and the nation. Therefore, it is imperative that organizations remain vigilant and that senior executives, leaders, and managers throughout the organization understand their responsibilities and are accountable for them. In addition to the responsibility to protect organizational assets from the threats that exist in today's environment, organizations have a responsibility to consider and manage the risks to protecting organizational assets and for managing risk. *It is in this regard that portfolio-based* **cyber-risk management (CRM)** *solutions (e.g., a mixture of self-defense solutions like anti-virus, commercial cyber-insurance, and self-insurance) are being invested in by organizations to manage the negative impact, both, tangible (e.g., monetary losses, devaluation of stock) and non-tangible (e.g., reputation, consumer trust), of the vulnerabilities exploited by cyber-hackers.*

Relating Our Research Contribution to MISs - In the event of a big cyber-attack on a retail company (e.g., the *Target* data breach attack), the latter incurs huge losses, arising from both

first-party sources (e.g., direct consumers) and third-party sources (e.g., clients of consumers). Though it might be feasible in the best case for CRM solutions to cover first-party losses with appropriate deductibles, third-party losses under high-impact cyber-attacks can add up from multiple sources to be too large for CRM solution providers to effectively manage. An inability to be effective in managing such large risk impacts has a huge negative impact on businesses (specifically the technologies, people, and processes) and society overall (well-being of people). The unresolved question of interest is: *can aggregate losses arising from both first-party and third-party sources after a big IoT-age cyber-attack jeopardize the CRM business so as to significantly change the means via which MISs function so as to enable the cost-effective, resilient, and cyber-risk free operation of IT-driven organizations?* We investigate this question and provide relevant insights.

2 WHY BE SKEPTICAL OF CYBER RE-INSURANCE SERVICES?

One of the key features of insurance as a business model is its ability to pool different types of risks, thereby reducing an underwriter's overall risk exposure. This is particularly true for a re-insurer, who is in a position to significantly diversify its risks by selling reinsurance to very different front-line insurers who specialize in different sectors (e.g., retail, pharmaceutical, manufacturing, etc.). This means that a re-insurer typically takes on a fraction of many different risks. There is also evidence suggesting that in cyber space, re-insurers increasingly act as co-insurers through the use of *proportional treaties* [Carter 2013], whereby the insurer and re-insurer split the cost of loss coverage on each policy proportionally according to the treaty/contract whenever loss occurs, rather than having the re-insurer only step in when the (aggregate) loss exceeds a (high) threshold specified by a non-proportional treaty.

In this section, we provide a mathematically intuitive explanation behind being skeptical of re-insurance services resolving the problem of successfully covering aggregate cyber-losses faced by a set of cyber-insurers, when an IT-driven liability-networked system is subject to a cyber-attack. In such situations, though the insurers of individual firms have the advantage of them diversifying their coverage to multiple other insurance firms of organizations they depend upon, there is a disadvantage angle to it as well that arises when a cyber-attack causes significant cascading losses in a network of service-liable organizations. *This implies the possibility of a single cyber-insurance company bearing the responsibility of covering an aggregate sum of individual organizational risks in a supply chain network.* Consequently, it is not certain here that a risk-averse re-insurance firm will always find it profitable (or even feasible) to cover aggregate losses for supply chain networks of IT-driven industries.

Having mentioned above about the summation of individual risks, it makes sense to investigate in the first place the impact that individual risk distributions might have on the aggregate risk after a cyber-attack. Traditional cyber-attacks often lead to organizational risk distributions that have short tails [Coburn et al. 2018]. On the contrary, modern cyber-attacks, fueled by the rise of large-scale IoT technology, are likely to generate organizational risk distributions that are heavy tailed in nature [Coburn et al. 2018]. In such settings, it is interesting to get an idea of (and compare) whether the resulting aggregate risk distribution (from multiple organizational nodes) at a re-insurer's end is favorable to provide coverage. We consider the *Normal* distribution as a representative of light-tail distributions, and the *Levy* and the *Cauchy* distributions as representative examples of heavy-tailed risk distributions that are stable,[5] i.e., a subclass of distributions whose

[5]A distribution is said to be stable if a linear combination of two independent random variables with this distribution has the same distribution, up to location and scale parameters. The Normal, Cauchy, and Levy distributions are the only stable distributions for which closed-form expressions exist and consequently help in tractable analyses.

left tails satisfy a Pareto law and exhibit power-law decay of the form $F(-x) \approx x^{-\alpha}$. Here $x, \alpha > 0$, and F is a **cumulative distribution function (cdf)** for a risk **random variable (r.v.)** X.

It is popular knowledge that for K i.i.d cyber-risk random variables $X_1, X_2, \ldots, X_K$ chosen from the standard normal $\mathcal{N}(\mu, \sigma^2)$, the resultant r.v. $\frac{\sum_{i=1}^{K} X_i}{K}$ is distributed with $\mathcal{N}(\mu, \sigma^2)$. The system implication of this r.v. in our article setting is a cyber-insurance company that outsources risk X_i to re-insurer i among the K cyber re-insurers. Thus, the risk spread of the popular **value-at-risk (VaR)** metric [Holton 2003], reflected through the spread parameter σ, grows as $\sqrt{\frac{1}{K}}$ of σ for a given location parameter μ, implying a *decrease* in VaR[6] spread on sum-averaging K risks. Thus, in this case it is better for a cyber-insurance company to re-allocate/spread the risks from its clients to re-insurers. Now consider a cyber-risk distribution that is Levy distributed [Forbes et al. 2011] with location parameter μ and spread parameter σ. The pdf is given by

$$\phi(x) = \begin{cases} \sqrt{\frac{\sigma}{2\pi}} e^{\frac{-\sigma}{2(\mu - x)}} (\mu - x)^{\frac{-3}{2}} & \text{if} x < \mu, \\ 0 \quad \text{if} x \geq \mu, \end{cases}$$

and the cdf is

$$F(x) = \begin{cases} \frac{2}{\sqrt{\pi}} \int_0^{\frac{-\sigma}{\sqrt{2(\mu - x)}}} e^{-t^2} dt & \text{if} x < \mu. \\ 1 \quad \text{if} x \geq \mu. \end{cases}$$

Let $L_{\mu,\sigma}$ be the class of r.v.s with the above Levy distributions. Thus, for K i.i.d cyber-risk random variables $X_1, X_2, \ldots, X_K$ chosen from $L_{\mu,\sigma}$, we get $\frac{\sum_{i=1}^{K} X_i}{K} \in L_{\mu, K\sigma}$. Therefore, contrary to the case of the Normal distribution, the value-at-risk spread σ in the case of the Levy distribution *increases* K-fold for a given μ, implying a K-fold increase in cyber-risk on sum-averaging K risks. Thus, in this case it might not be beneficial for a re-insurance company to accept the multiple risks from its clients. As another example, take the Cauchy distribution, whose pdf for a given location parameter μ and scale parameter σ is given by

$$\phi(x) = \frac{1}{\pi\sigma} \frac{1}{1 + \left(\frac{(x-\mu)^2}{\sigma^2}\right)},$$

where $\sigma, X \in S_{\mu,\sigma}$—the set of r.v.s with Cauchy distribution having the corresponding location and spread parameters. The cdf of X is given by

$$F(x) = \frac{1}{2} + \frac{1}{\pi} \tan^{-1}\left(\frac{x - \mu}{\sigma}\right).$$

Thus, for K i.i.d cyber-risk random variables $X_1, X_2, \ldots, X_K$ chosen from $S_{\mu,\sigma}$, we get $\frac{\sum_{i=1}^{K} X_i}{K} \in S_{\mu,\sigma}$. Therefore, contrary again to the case of the Normal distribution, the value-at-risk spread σ in the case of the Cauchy distribution does not decrease for a given μ, implying neither an increase nor a decrease in cyber-risk on sum-averaging K risks. The Cauchy case is this intermediate between the Levy case and the case with Normal distributions.

[6]We use the VaR notion of cyber-risk measure due to the fact that heavy-tailed distributions like the Levy and Cauchy distributions do not have finite first- or second-order moments [Forbes et al. 2011]—hence functions of expected measures of cyber-risk variables are undefined. One could well argue the use of the popular expected shortfall, i.e., CVaR, a cyber-risk measure that is coherent and is defined as the average of the worst losses of a portfolio; however, this metric requires existence of the statistical first moments of cyber-risks to be finite, which may not be true of catastrophic cyber-risks. Thus, the feasibility connotations with respect to the VaR metric would coincide with that obtained with respect to the CVaR metric. In addition, it is not difficult to see from Acerbi [2002] and Cotter and Dowd [2006] that the assumptions close to the existence of the means of the cyber-risks in consideration are also required for applications of coherent spectral measures of cyber-risk that generalize expected shortfall.

Table 1. Table of Important Notations

X	random cyber-risk
k	coverage liability limit
s	number of risk sharing cyber-insurers
j	number of risks to be shared by the (re-)insurers
$V(\cdot)$	effective expected outcome for a cyber-insurer
$U_{j,s}$	expected utility of s cyber-insurers sharing j risks equally
M	maximum number of cyber-insurers in the re-insurance market
N	maximum number of cyber-risks to be shared
$S_{\nu,\sigma}$	Cauchy random variable, location parameter ν, scale parameter σ
$X_i \in S_{\nu,\sigma}$	i.i.d. Cauchy random variables
μ	cyber-insurance premium in terms of number of loss units
α	risk-averse parameter for a cyber-insurer
q	probability of a cyber-attack

Intuition-Driven Practical Insight - It is somewhat clear that light-tailed distributions might pose less VaR to cyber re-insurers when compared to heavy-tailed distributions. Even for the case when $c_i \in \mathbb{R}_+ | \sum_{i=1}^{K} c_i = 1$, instead of being a uniform $\frac{1}{K}$, for each $i \in \{1, \ldots, K\}$, we will have $\sigma = (\sum_{i=1}^{K} (c_i \sigma_i)^{\frac{1}{2}})^2$ in case of the Levy distribution, and $\sigma = \sum_{i=1}^{K} c_i \sigma_i$ in case of the Cauchy distribution. In both these distributional scenarios, the VaR to re-insurers is more than in the case when some $c_i = 1$, for a given i, and $c_j = 0$ for all $j \neq i$ (follows from the application of results in *majorization theory* [Marshall et al. 1979] (see appendix for a brief overview). This puts weight on our skepticism that cyber re-insurance services may not be profitable in the case when individual insurers with liability limits are faced to cover heavy-tailed cyber-risks. Note that our skepticism also extends to scenarios where cyber-risk distributional supports are bounded (e.g., under limited risk liabilities as mentioned in subsequent sections) for which an expected utility analysis on first moments can be conducted.

3 A FORMAL MODEL OF A CYBER RE-INSURANCE MARKET

In this section, we present a game-theory-driven strategic market model, based on developments in Ibragimov et al. [2009], to characterize the (in)-effectiveness of cyber-re-insurance markets for effective aggregate cyber-risk coverage of catastrophic cyber-risks. To this end, we first provide the setup of our model and its qualitative overview. We then follow it up with the description of a dynamic game-induced cyber re-insurance market model and an equilibrium existential analysis. A list of important notations used in our paper is stated in Table 1.

3.1 Model Setup

We assume that there are M cyber-insurance companies in operation in a market, and each of them is an expected utility optimizer with identical strictly concave utility functions u. Each cyber-insurer takes up a limited liability on consumer cyber-risk coverage, as in real practice. This is rational on the part of the cyber-insurer for two reasons: (1) for heavy-tailed cyber-risk distributions, there is a non-zero probability that a risk-averse cyber-insurer might go default on covering an exorbitant amount of cyber-risk due to a successful cyber-attack on its insured firms, and (2) without a liability limit, one cannot do an expected value analysis on cyber-insurer utility functions, simply because the first moments are undefined for catastrophic heavy-tailed distributions.[7]

[7]Note that not all heavy-tailed distributions represent catastrophic cyber-risks [Zolotarev 1986], and for such non-catastrophic distributions, the mean of the distribution is finite.

Without a loss of generality, we study the strategic feasibility of aggregate cyber-risk coverage with curtailed coverage limits via an expected utility theoretic analysis.[8] The assumption of limited liability is modeled by cyber-insurers being liable to cover losses only up to a certain amount k. If losses exceed k, an insurer defaults on any additional loss beyond k and might resort to re-insurance services for covering the additional loss. Thus, if a cyber-insurer takes on a random risk X, the effective outcome (before opting for cyber re-insurance services) for the insurer once X is realized is

$$V(x) = \begin{cases} X & \text{if} X \geq -k, \\ -k & \text{if} X < -k. \end{cases} \tag{1}$$

In the special case when there is no limited liability, i.e., when $k = \infty$, we have $V(X) = X$ for all X. If $k < \infty$, u is defined only on $[-k, \infty]$, and without loss of generality, $u(-k) = 0$. Considering the assumption that cyber-risks $X_1, \ldots, X_M$ are i.i.d.,[9] we wish to study the expected utility of s cyber-insurance agents who share j risks equally. To this end, we define the random variable $z_{j,s} = \frac{\sum_{i=1}^{j} X_i}{s}$, with cdf $F_{j,s}$. The expected utility of such risk sharing is given by

$$U_{j,s} = \mathbb{E}(V(z_{j,s})) = \int_{-k}^{\infty} u(x) dF_{j,s}(x). \tag{2}$$

As mentioned above, the cyber-insurance firms are assumed to be risk averse, simply because executives with major financial and human capital investments in their own firms wish to avoid risky situations. The value of firms in this case is a concave transformation of the payoffs and is effectively identical to an expected utility setup [Froot et al. 1993]. We assume that each cyber-insurer can bear risks for a maximum of N asset lines[10] spread across its client organizations. Thus, we have $1 \leq s \leq M$ and $1 \leq j \leq Ns$. We also assume that each X_i is non-divisible. This assumption makes sense in practice because not every cyber-insurer will insure cyber-risks from all possible asset lines from all possible geographical locations (e.g., could be constrained by geographical policies). It is then evident from Section 3.1 that when X_is are normally distributed, $U_{j,s}$ is an increasing function in both s and j, and we can expect the cyber-reinsurance market to work well under light-tailed cyber-risks and insurance to be offered for the maximal number of cyber-risks NM as each cyber-insurer can bear a capacity of N risks from each of the $M - 1$ cyber-insurers. With respect to heavy-tailed risks, we assume i.i.d Bernoulli-Cauchy distributed cyber-risks $\tilde{X}_i$ of the following form, for the purpose of analytical tractability, primarily due to their *stable* nature:

$$\tilde{X}_i = \begin{cases} \mu & \text{with probability } 1 - q, \\ X_i,\ X_i \in S_{\nu,\sigma} & \text{with probability } q, \end{cases}$$

where $X_i \in S_{\nu,\sigma}$ are i.i.d. Cauchy random variables with location parameter ν and scale parameter σ. Alternatively, the r.v.s $\tilde{X}_i$ are "mixtures" of degenerate and Cauchy r.v.s and can be written as

$$\tilde{X}_i = \mu(1 - \epsilon_i) + X_i \epsilon_i = \mu + (\nu - \mu)\epsilon_i + \sigma Y_i \epsilon_i, \tag{3}$$

[8]The orthogonal problem of deriving the statistical infeasibility of aggregating catastrophic cyber-risks with limits on coverage liability and using an expected utility theoretic analysis is shown via Theorem 7.4 of the appendix. This result establishes the equivalence of the statistical infeasibility of aggregate catastrophic cyber-risk coverage achieved via both a VaR-centric analysis (not the focus of the article) and an expected utility-driven analysis (the focus in this article).

[9]Note here that we are considering catastrophically heavy-tailed distributions that are i.i.d. The popular notion (also mentioned in Section 1) is that non-catastrophic cyber-risks that are heavy tailed and dependent are a cause of concern for cyber-insurance coverage markets. In this work we will show that contrary to popular intuition, even i.i.d. catastrophic cyber-risks are a cause of concern for commercial insurance coverage.

[10]Each asset line might be a part of a coverage type. For example, the category of first-party coverage could include business interruption, restoration, and crisis communications as different asset lines, whereas the category of third-party coverage could include data breaches, network interruption, and notification expenses as examples of asset lines.

where ϵ_i are i.i.d. non-negative Bernoulli r.v.s with $P(\epsilon_i = 0) = 1 - q, P(\epsilon_i = 1) = q$, and $Y_i \, \epsilon \, S_{0,1}$ are i.i.d. Cauchy symmetric r.v.s with scale parameter $\sigma = 1$, that are independent of ϵ_is. We say that $\tilde{X}_i \in \tilde{S}^q_{\mu,\nu,\sigma}$, where μ is a reflection of the premium a cyber-insurance provider collects to insure against loss events that occur with probability q. We assume the utility function of a risk-averse cyber-insurer to be of the following form:

$$u(x) = (x + k)^\alpha, \; \alpha \in (0, 1),$$

which is the power utility function. x, being a risk variable, is a **Von-Neumann Morgentern (VNM)** utility function. α is the degree of risk aversion of the cyber-insurer. Note that under the limited liability assumption, the expected utility for Bernoulli-Cauchy risks always exists.

3.2 A Qualitative Overview of Our Cyber Re-Insurance Market Model

We consider a market setting consisting of primary (traditional) strategic cyber-insurance providers selling insurance against cyber-risks. Either some or all of these insurers can further pool together to form a cyber re-insurance market. The setup is a two-stage game that captures the intuitive idea that primary cyber-insurance services need to be offered before re-insurance can be pooled. The decision of an individual cyber-insurer whether to offer primary cyber-insurance solutions will be based on its belief of how well-functioning cyber-reinsurance markets will be in the future. In the first stage of the game, strategic cyber-insurers simultaneously choose whether to offer primary insurance against a set of i.i.d. cyber-risks. This assumption makes sense because the decision to offer primary insurance is the least conservative in the presence of i.i.d. cyber-risks. A negative decision here implies a negative decision for non i.i.d. risks as well, and one of the goals in this article is to study the most conservative conditions under which a (re-)insurance market can fail. In the second stage of the game, the cyber re-insurance market is formed and each cyber-insurer decides whether to participate or not. Cyber-insurers who decide not to provide primary insurance are allowed to participate in the re-insurance market. Finally, all risks of insuring entities participating in the re-insurance market are pooled and outcomes are realized and shared equally among the participating re-insurers. We assume the latter primarily for tractable simplicity, and secondarily because of the fact that the cost to garner information that would result in 'proportionally' fair re-insurer allocations might be expensive enough to result in market failure, and to avoid such scenarios a regulatory intermediary might settle on equal sharing behavior.

3.3 Formal Market Model

Stage 1 of Market Game - We assume $M \geq 2$ cyber-insurance providers, indexed by $1 \leq m \leq M$. There is a set of i.i.d. cyber-risks $\mathcal{X}$, where each has a distribution (cdf) $F(x)$ for $x \in \mathcal{X}$. Each cyber-insurer chooses to take on a specific number of risks, $n_m \in \{0, 1, \ldots, N\}$, on different asset lines, where N denotes the maximum insurance capacity forming a portfolio of cyber-risks, $p_m \in \mathcal{P}$, where $p_m = \sum_{i=1}^{n_m} X_i$, and $X_i \in \mathcal{X}$. This is the first stage of the market. The cyber-risks are assumed to be atomic (indivisible) in nature and each risk can be chosen by at most one cyber-insurer. As previously mentioned, this assumption reflects the situation that insurers may be tied by their corporate policies to cover certain, but not all, lines of cyber-risk—thus we disallow divisible risks in our model. We also assume that there are enough cyber-risks available to exhaust capacity, i.e., $|\mathcal{X}| = NM$. In addition, we assume there are enough cyber-risks for a single cyber-insurer not to be able to handle all of them alone. Here, $|\mathcal{X}|$ denotes the cardinality of $\mathcal{X}$. As the cyber-risks are i.i.d., only the distributional assumptions of the risks matter, so we do not concern ourselves with which cyber-insurance provider chooses which risk to cover as long as it falls within one's allowable set. The portfolio p_m is thus characterized completely by the number of risks n_m. The total number of cyber-risks insured is $\bar{N} = \sum_m n_m$. Each cyber-insurer has the liability to cover cyber-losses up to

an amount $k \in [0, \infty]$. When cyber-losses exceed k for an insurer, it defaults and first pays off k, and then it approaches a third party, either a re-insurer or the government, to cover the excess losses. The effective outcome under limited cyber-risk liability for an insurer taking on risk z_m is $V(z_m)$, where $V(\cdot)$ is defined as in Equation (1). Each cyber-insurer is assumed (for tractable simplicity) to have identical expected utility over cyber-risks, $U_m(z_m) = \mathbb{E}u(V(z_m))$, where u is defined and continuous on $[-k, \infty]$, is strictly concave, is twice continuously differentiable on $(-k, \infty)$, and, if $k < \infty$, satisfies $u(-k) = 0$. The outcome of the first stage is summarized by $p = (p_1, \ldots, p_M) \in \mathcal{P} = \Pi_{m=1}^{M} \mathcal{P}_m$.

Stage 2 of Market Game - In the second stage of the game, the cyber-re-insurance market is formed. In this stage, the cyber-insurers are assumed to have perfect knowledge about p, the outcome of the first stage of the game. Each cyber-insurer, $1 \leq m \leq M$, sequentially decides whether to participate in the re-insurance market or not, via the following rationale: First, cyber-insurer 1 decides whether to participate in the market. This is represented through a binary variable $q_1 \in \{0, 1\}$, where $q_1 = 1$ denotes that cyber-insurer 1 participates in the cyber re-insurance market and $q_1 = 0$, otherwise. Then cyber-insurer 2 decides whether to participate, observing insurer 1's decision. This is repeated until all M cyber-insurers have decided. We assume that previous cyber-insurer decisions are perfectly observable. If a cyber-insurer is indifferent between participating and not participating, it will not participate in the pooling cyber re-insurance market. The payoff to not participating is denoted as $U_{n_m,1}$, and the payoff to participating is denoted as $U_{R,t}$, where $t = \sum_m q_m n_m$ denotes the number of participating cyber-insurers and $R = \sum_m q_m n_m$. Cyber-insurers who agree to be part of the re-insurance market pool all their primary cyber-insurance in the re-insurance market, i.e., amount $q_m p_m$. The total pooled risk is therefore $P = \sum_m q_m p_m$ and the number of risks $R = \sum_m q_m n_m \in \{0, \ldots, NM\}$. The outcome of the stage 2 subgame is summarized by $q = (q_1, \ldots, q_M) \in \{0, 1\}^M$, and the outcome of the overall game $\mathcal{G}$, named as the **re-insurance market game (RIMG)**, is completely characterized by (p, q). Note that RIMG is also characterized completely by the tuple $\mathcal{G} = (u, F, k, N, M)$.

3.4 The Existence of Market Equilibria in the RIMG

In this section, we investigate the existence and uniqueness aspects of market Nash equilibria for both stages of the RIMG.

The **second stage** of the RIMG is an M-step sequential game with perfect information, and so it is straightforward to notice that the unique subgame perfect Nash equilibrium arises as an application of the standard backward induction technique in game theory [Maschler Michael 2013], the rationale being that Zermelo's theorem [Maschler Michael 2013] immediately applies and implies that for each realization of the first stage of the RIMG, $p = (p_1, p_2, \ldots, p_m)$, there is a unique subgame perfect Nash equilibrium satisfying the laziness assumption (if cyber-insurers are indifferent between participating and non-participating, i.e., $U_{n_m,1} = U_{R,t}$, then they do not participate) to the participation stage of the RIMG, $q \in \{0, 1\}^M$. We can therefore define the equilibrium mapping $\varepsilon : \mathcal{P}^M \to \{0, 1\}^M$, with $q = \varepsilon(p)$. Without loss of generality, we can assume that in the first stage of the RIMG, cyber-insurers base their strategies on this equilibrium mapping. This reduces the strategy space significantly without having any effect on the subgame perfect Nash equilibrium outcome. For elements $p \in \mathcal{P}$, we define the **first-stage** actions of all cyber-insurers except insurer m as

$$p_{-m} = (p_1, \ldots, p_{m-1}, p_{m+1}, \ldots, p_M) \in \Pi_{m' \neq m} \mathcal{P}_{m'} = \mathcal{P}_{-m}.$$

A strategy for cyber-insurer m consists of a pair $A = (p_m, \eta_m) \in \mathcal{P}_m \times \{0, 1\}^{\mathcal{P}_{-m}}$, where p_m is the chosen portfolio of cyber-insurance and $\eta_m : \mathcal{P}_{-m} \to \{0, 1\}$ is the participation choice, depending on the realization of the first stage of the RIMG. Cyber-insurer m has a belief set about the other

insurers' first-stage actions, $B_m = p_{-m} \in \mathcal{P}_{-m}$. Insurer m's strategy, $A_m = (p_m, q_m)$, conditioned on belief set $B_m = p_{-m}$, is said to be consistent if $\eta_m(p_{-m}) = \varepsilon(\tilde{p}))_m$, where

$$\tilde{p} = ((p_{-m})_1, \ldots (p_{-m})_{m-1}, p_m, (p_{-m})_{m+1}, \ldots, (p_{-m})_M), \tag{4}$$

where we use the notation $(x_i)_i$ for the ith element of the ordered set x. Rational cyber-insurers will consider only consistent strategies when compared to non-consistent ones as they are sub-optimal in the participation phase of the RIMG. The inferred outcome of a consistent strategy, $A_m = (p_m, \eta_m)$, conditioned on a belief set B_m is

$$z_m(p_m|B_m) = \begin{cases} p_m & \text{if } \eta_m(p_{-m}) = 0 \\ \frac{P}{t} & \text{if } \eta_m(p_{-m}) = 1, \end{cases}$$

where $\tilde{q} = \varepsilon(\tilde{p})$, $t = \sum_{m'}(\tilde{q})_{m'}$, $P = \sum_{m'}(\tilde{p})_{m'}(\tilde{q})_{m'}$, and $\tilde{p}$ is as defined in Equation (4).

An M-tuple of strategies, $(A_1, \ldots, A_M)$ and belief sets $(B_1, \ldots\ldots, B_M)$, where $A_m = (p_m, \eta_m)$ and $B_m = p_{-m}$, defines a RMNE of the RIMG, if

(1) Consistent Strategies: For each cyber-insurer m, A_m is consistent conditioned on belief set B_m.
(2) Maximized Strategies: For each cyber-insurer m, $p_m \in \text{argmax}_{p' \in P} U_m(z_m(p'|B_m))$.
(3) Consistent Beliefs: For each cyber-insurer m, for all $m' \neq m : (p_{-m})_{m'} = p_{m'}$.

The RMNE outcome is summarized by $p = (p_1, p_2, \ldots, p_M)$ and $q = (\eta_1(p_{-1}), \eta_2(p_{-2}, \ldots, \eta_M(p_{-M})$.

4 MARKET EQUILIBRIUM ANALYSIS WITH PRACTICAL IMPLICATIONS

In this section, we derive results for our application setting, built upon constructs in Ibragimov et al. [2009] that relate multiple RMNE's of the RIMG to various conditions on the $U_{j,s}$-defined in Equation (2). We provide practical implications of these results on the cyber-(re)insurance market, and lay down the proofs of all the results in Section 6.

Theorem 4.1. *If $U_{j,s} < U_{0,1}$ for all $j \in \{1, \ldots, N\}$ and all $s \in \{1, \ldots, M\}$, there exists an RMNE of the RIMG in which no cyber-insurance is offered, and the equilibrium is unique.*

Practical Implications - The theorem states that the condition that expected utility $U_{0,1}$ of not sharing any risk, i.e., the primary cyber-insurance case, will be greater than that, i.e., $U_{j,1}$, of sharing $j \in \{1, \ldots N\}$ risks by any single cyber-insurer *discourages a primary cyber-insurance market formation in the first place.* In the age of cascading liability cyber-risks, this makes sense because there is a considerable likelihood that a single cyber-insurer, in the event of heavy-tailed cyber-risks, would have to resort (cyber-risks may be too large) to other insurers for covering client losses, and if the latter do not find it incentive compatible to provide re-insurance services, it is not economically incentive compatible either for a cyber-insurer to be part of a primary coverage market as it might not be able to satisfy coverage QoS of its clients. Moreover, even if there exists a cyber re-insurance market, there is no way to increase the expected utility by cyber-risk sharing if only one cyber-insurer contributes risk to the cyber re-insurance market. As a problem-alleviating step, regulators should take note of such situations and intermediate appropriately (see Section 7.2 for more details) in the premium standardization process so that an increasing number of cyber-insurers are incentivized to pool cyber-risk.

Theorem 4.2. *If $U_{j,s} > U_{0,1}$ for some $j \in \{1, \ldots, N\}$ and $s \in \{1, \ldots, M\}$, then there does not exist an RMNE of the RIMG in which no cyber-insurance is offered.*

Practical Implications - The theorem states that under condition C1 being not satisfied, unlike in the previous case, it *does not discourage* a single cyber-insurer towards primary cyber-insurance

market formation. However, the condition does not *necessarily guarantee* the existence of a market equilibrium encouraging a cyber-insurer towards primary cyber-insurance market formation.

Theorem 4.3. *The RIMG leads to the existence of an RMNE if the following conditions hold:*

(1) $U_{NM,M} > U_{j,1}$ *for all* $j \in \{0, \ldots, N\}$.
(2) $U_{NM,M} > U_{j,M}$ *for all* $j \in \{N(M-1), \ldots, NM-1$.

Under such an RMNE, cyber-insurance against all risks in $\mathcal{X}$ *is offered, and all the cyber-risk insured is pooled in the cyber re-insurance market.*

Practical Implications - The theorem extends the benefits obtained on conditions satisfied in Theorem 4.2 to the case where there is a guaranteed RMNE that allows providing cyber-insurance against all cyber-risks, as an economically incentive-compatible solution. More importantly, a market for re-insurance services is also deemed economically incentive compatible. Conditions 1 and 2 in the theorem act as sufficient conditions for the successful existence of cyber re-insurance markets. The conditions imply greater expected utility on pooling all re-insurable cyber-risks together as a team when compared to that of handling at least one re-insurable cyber-risk by a single cyber-insurer. Similar to that in Theorem 4.1, regulators will have a big role to play to effectively regulate cyber-insurance premiums that satisfy such conditions.

Theorem 4.4. *There is a possibility of two starkly contrasting RMNEs:* $RMNE_1$ *and* $RMNE_2$*, existing in the RIMG where at* $RMNE_1$ *all cyber-risk is pooled in the cyber re-insurance market, and at* $RMNE_2$ *no primary cyber-insurance (subsequently cyber re-insurance) service is offered, despite a large cyber-risk-bearing insurance capacity available in the market. Moreover, this multiple equilibrium situation exists if there is an* M_0 *for all* $M \geq M_0$ *for the RIMG* $\mathcal{G} = (u, F, k, N, M)$.

Practical Implications - The possibility of the existence of two starkly contrasting RMNEs is a disadvantage from a social welfare perspective, simply because there exists a situation as a function of M_0 where despite the pooling capacity available to re-insure cyber-risks, the strategic cyber-insurance firms do not find it economically incentive compatible to participate in the re-insurance, and worse, even primary insurance markets, resulting in a market failure and sub-optimal welfare state. However, one good thing is that this situation cannot arise when cyber-risk distributions have a finite mean and variance, which is the case for the more probable non-catastrophic cyber events. The disadvantageous two-RMNE situation is similar to the outcome of the popular Prisoner's Dilemma game in non-cooperative game theory and calls for regulatory action to effectively coordinate actions between the cyber-insurers to reach the equilibrium that improves/maximizes social welfare.

Theorem 4.5. *If (i)* $k = \infty$ *and the cyber-risks* $X \in \mathcal{X}$ *have finite second moments,* $\mathbb{E}(X^2) < \infty$*, or (ii)* $k < \infty$ *and the cyber-risks* $X \in \mathcal{X}$ *have* $\mathbb{E}(X^2) < \infty$*, and* $\mathbb{E}(X) \neq 0$*, or (iii)* $k < \infty$ *and the cyber-risks* $X \in \mathcal{X}$ *have* $\mathbb{E}(X^{2+\epsilon}) < \infty$ *for some arbitrarily small* ϵ*, and* $\mathbb{E}(X) = 0$*, then there exists no* $M \geq M_0$ *for the RIMG* $\mathcal{G} = (u, F, k, N, M)$*, for which there is a possibility of two starkly contrasting RMNEs:* $RMNE_1$ *and* $RMNE_2$*, existing in the RIMG where at* $RMNE_1$ *all cyber-risk is pooled in the cyber re-insurance market, and at* $RMNE_2$ *no primary cyber-insurance (subsequently cyber re-insurance) service is offered, despite a large cyber-risk-bearing insurance capacity available in the market.*

Practical Implications - The implications are similar to those of Theorem 4.4 that talk about the ramifications of two-equilibria market states, except that here we enforce a stricter necessary condition on the cyber-risk variables for the disadvantageous two starkly contrasting market equilibria state to arise. More specifically, (1) when there are no limits on cyber-insurer liability, cyber-risk distributions are required to have infinite second moments for two-equilibria states to

arise, whereas (2) under the limited liability regime, cyber-risk distributions need to have both infinite means and infinite variances for two-equilibria market states to arise. *Both these conditions indicate a propensity towards a successful presence of re-insurance markets for non-heavy/not-so-heavy-tailed cyber risks, a trait not shown by low-probability-occurring catastrophic cyber-events.* In summary, contrasting two-equilibria market states for any $M \geq M_0$ can arise *if and only if* cyber-risk distributions have catastrophe reflecting heavy tails.

Thus far, we analyzed the feasibility of cyber re-insurance under catastrophic cyber-risks under a curtailed risk distribution model. Based on the theory of stable distributions and majorization theory (see Appendix A), we have the following result (courtesy of Ibragimov [2009]), showing that cyber re-insurance *may not be a viable market proposition* under the VaR metric-induced framework for providing insurance against aggregate cyber-risks for non-curtailed risks that hit individual organizations and having infinite first moments and unbounded distribution support. This result generalizes the fact that one must be cautious of implementing cyber re-insurance markets under catastrophic events irrespective of whether (1) the cyber-risk distributions have finite or infinite first moments and (2) the cyber-risk distributions have bounded or unbounded support.

Theorem 4.6. *Let $q \in (0, \frac{1}{2})$ and let $X_i, i = 1, \ldots, n$, be i.i.d. cyber-risks such that $X_i \sim CS(r), i = 1, \ldots, n$. Then*

(1) $VaR_q(Z_v) > VaR_q(Z_w)$ if $v \prec w$ and v is not a permutation of w (in other words, the function $\psi(w, q) = VaR_q(Z_w)$ is strictly Schur-concave in $w \in \mathrm{R}^n_+$).
(2) In particular, $VaR_q(Z_w) < VaR_q(Z_w) < VaR_q(Z_w)$ for all $q \in (0, \frac{1}{2})$ and all weights $w \in \mathcal{I}_n$ such that $w \neq \underline{w}$ and w is not a permutation of $\bar{w}$.

Practical Implications - The theorem implies corresponding results on majorization properties of the tail probabilities $\psi(w, x) = P(\sum_{i=1}^n w_i X_i > x), x > 0$, of linear combinations of heavy-tailed r.v.s $X_1, \ldots, X_n$: these implications generalize the results in the seminal work by Proschan [1965], who showed that the tail probabilities $\psi(w, x)$ are Schur-convex in $w = (w_1, \ldots, w_n) \in R^n_+$ for all $x > 0$ for i.i.d. r.v.s $X_i \sim \mathcal{LC}, i = 1, \ldots, n$. Schur-convexity of $\psi(w, x)$ for $X_i \sim \mathcal{LC}$ implies that the value at-risk comparisons in the theorem hold for i.i.d. log-concavely distributed cyber-risks. Now let us consider the portfolio value at risk dealt with in the theorem in the borderline case $\alpha = 1$, which corresponds to i.i.d. cyber-risks $X_1, \ldots, X_n$ with a symmetric Cauchy distribution $S_1(\sigma, 0, 0)$. These distributions are exactly at the dividing boundary between the class $\overline{\mathcal{CSLC}}$ and the class $\underline{\mathrm{CS}}(1)$. With $\alpha = 1$, we get that, for all $w = (w_1, \ldots, w_n) \in \mathcal{I}_n, Z_w = \sum_{i=1}^n w_i X_i =_d X_1$. Consequently, for all $q \in (0, 1)$, the value at risk $VaR_q(Z_w) = VaR_q(X_1)$ is independent of w and is the same for all portfolios of cyber-risks X_i with weights $w \in \mathcal{I}_n, i = 1, \ldots, n$. Thus, in such a case, diversification of a portfolio has no effect on riskiness of its return. Similarly, for general weights $w = (w_1, \ldots, w_n) \in \mathrm{R}^n_+$, and $\alpha = 1$ implies $Z_w = \sum_{i=1}^n w_i X_i =_d (\sum_{i=1}^n w_i) X_1$. Thus, the value at risk $VaR_q(Z_w) = (\sum_{i=1}^n w_i) VaR_q(X_1)$ is independent of w so long as $\sum_{i=1}^n w_i$ is fixed. Consequently, $VaR_q(Z_w)$ is both Schur-convex and Schur-concave in $w \in \mathrm{R}^n_+$ for i.i.d. risks $X_i \sim S_\sigma(\sigma, 0, 0)$ that have symmetric Cauchy distributions with $\alpha = 1$.

We also analyze aggregate cyber-risk coverage feasibility for risk managers who are quite risk averse not to comply with the VaR risk measure. We already know from a well-celebrated utility-theoretic result due to Samuelson that for general (non-heavy-tailed) cyber-risks with bounded statistical support, aggregate risk coverage is always feasible for a risk manager. To this end, we extend this utility-theoretic result for curtailed catastrophic cyber-risks having heavy tails via the following result.

Theorem 4.7. *Let $n \geq 2$. Then there exists a t_0, such that for any $t \geq t_0$, there is an admissible utility function u and $a > 0$, such that any cyber-risk manager with utility function, v, where v is a*

t-convex regularization of u, will have

$$Ev\left(Y_1(a)\right) > Ev\left(\bar{Y}_n(a)\right).$$

Practical Implications - The theorem simply implies that the feasibility of covering aggregate cyber-risk crucially depends on the tail properties of the expected utility function and that if managers' utility function at any point in the domain of large negative outcomes becomes convex, then the non-feasibility of covering aggregate heavy-tailed cyber-risks may continue to hold. *On a practical note, this provides additional support for our view that the theory of unbounded catastrophic cyber-risk distributions may provide a good approximation for markets with a limited number of bounded catastrophic cyber-risks.* In light of this theorem, it is clear that in situations when we can assume that cyber-risk managers' utilities are strictly concave in the whole (efficient) support of distributional outcomes, we expect tradition feasibility (for non-catastrophic cyber-risks) results to hold whenever cyber-risks are bounded. However, in situations when the number of cyber-risks is not large compared with the number of liabilities, and if cyber-risk manager utility is non-concave for large negative outcomes, then it may be optimal for the latter to not participate in aggregate cyber-risk coverage even with bounded risks.

5 EXPERIMENTAL EVALUATION

In this section we validate our theory proposed in the article by first showing experimentally, via details from a corporate case study[11] on a cyber-attack launched on a real-world US-based cloud provider, the heavy-tailed possibility of cyber-risk. Once we establish the chances of heavy-tailed cyber-risk, we run synthetic experiments on heterogeneous cyber-insurers and compare and comment on the results with those obtained from theory in a homogeneous cyber-insurance setting. *Our efforts serve as a mere conservative projection of larger-scale IoT-driven supply chain service networks.*

5.1 The Real-World Case Setting

Organizational Knowhow - We deal with the case of a US-headquartered multi-national **cloud service provider (CSP)**, *A*, with a market share on the order of major public CSPs headquartered in the United States. CSP A operates multiple regions and many data centers around the world,

[11] In an orthogonal study [Welburn and Strong 2019], the authors *provide non-statistical point estimates (via a model) of aggregate per-day monetary cyber-loss impacts due to cyber-attacks on three large and likely targeted firms.* The firms chosen were (1) *AT&T,* a large telecommunications firm; (2) *Cisco Systems,* a hardware firm; (3) *JP Morgan Chase,* a retail banking firm; and (4) *Visa,* a point-of-sale firm, for their diversity, size, and potential exposure to cyber risks. To analyze the impact of atypically large yet still fathomable cyber-incidents, the authors assumed a stoppage of all operations and services (subject to resilience) of each company lasting 1 day. The authors projected (as point estimates) "upper bounds" of potential losses associated with each incident rather than the statistical expectation of losses, where actual losses could be smaller (in fact, due to cyber risk management strategies of individual firms and the resilience across supply chains). For Visa, the direct cost of USD 33 million incurred by the 1-day outage could lead to a further USD 77 million in upstream losses. The results of attacks on the other firms were similar: *upstream* supply chain losses to Cisco, JP Morgan Chase, AT&T, and Ford were estimated to be bounded above by USD 209, USD 145, USD 700, and USD 706 million, respectively. AT&T was projected to have the largest *downstream* impact in its supply chain, potentially USD 20 billion in damages in addition to the direct and upstream impacts. Notably, with Visa being the far smaller (in annual revenue base) firm compared to AT&T in our scenario, it had the potential for the second largest total impact from cascading failures downstream, reaching USD 13 billion, due to its huge client base. Furthermore, in stark contrast to AT&T, a firm of similar size, Ford's potential downstream costs were estimated to be much lower at USD 6.4 billion. Cisco and JP Morgan Chase were estimated to have downstream impacts of nearly USD 6 billion and USD 4 billion, respectively. *These point estimate calculations reveal the staggering potential impacts of systemic cyber risk, with potential per-day losses incurred by certain businesses that might shoot up to USD 20 billion, representing roughly 0.1% of US GDP.* The same scale of per-day average downstream impacts were point-estimated for *Maersk* on real data obtained after the *NotPetya* cyber-attack.

with five regional hubs to serve its customers in the United States. It has other hubs in many other parts of the world to serve its international client base. CSP A employs a strong team of technical security specialists who develop tools and detection systems for potential malware and devise contingency plans and response protocols for a wide range of technical issues.

The Cyber-Attack Scenario - An area of security importance for CSP A is the space of potential vulnerabilities in their **Routing Information Protocols (RIPs)**, the controlling system for connecting customers to the servers in the data centers. They have a testing lab that is disconnected from the main network but is periodically reconnected to enable live tests of new security systems to be performed. Following one of these routine reconnections, problems started occurring with routers in the regional hub that controls the active data centers. A few initial failures and loss of computing capacity meant that the operators opened reserve channels to draw capacity from other regional hubs. The security team figured out that a malware—a binary worm—had somehow found its way into the operational networks of the data center. The binary worm makes copies of itself if it is not controlled by another software twin. The worm changed the Routing Information Protocols in the routers within the regional hub and its satellite data centers. The address book to the regional hub was effectively being erased, component by component. There are around 2,000 routers in the regional hub; some are core routers to direct the primary traffic flows, while others are edge routers that connect to customers.

Attack Spread and Single Regional Center Impact - The worm was virulent, rapidly self-replicating, and destructive. When it reconfigured the RIP of a router, the router could not be repaired through remote re-programming. It required a manual process of skilled operators to find the router in the racks of the server farms and spend several hours re-programming the firmware to enable the router to be brought back online and to resume its function. The only saving grace was that the worm was not disguised; i.e., it could be discovered and deleted. But the worm had infected the system extensively and the deletion of instances of the worm led to re-infection from unfound copies. When the worm attacked a router, the addresses controlled by that router were lost. The machines, networks, and data were unaffected and continued to run, but their users could no longer access them. Each affected customer found that their connection to their cloud service provided by CSP A no longer operated. Within an hour, over 5,000 companies found that their cloud service had failed. Around 1,000 more of CSP A's premium rate companies had switched their operations from the affected regional center to their alternate deployment on one of the other four of CSP A's centers.

Impact of Spread to Other Regional Centers - The Routing Information Protocol operated by CSP A allows for a limited load-balancing and software transmission between their different regional hubs. These interactions are routinely security screened to prevent malware transmission between centers. However, the emergency protocols for load balancing and transfer of capacity from one hub to another allowed reduced screening on high volumes of data traffic. It became rapidly apparent that the early-stage attempts to compensate for the loss of router capacity in the affected regional center had allowed the binary worm to spread to two other regional centers. As capacity began to be lost from these three infected hubs, CSP A's operational control personnel rapidly isolated all of their other regional hubs and prevented any interaction between them to avoid infection spreading to more centers and potentially affecting their worldwide operations of many regional hubs in other countries. The three infected hubs suffered a complete general system failure for all of their connected customers within an hour of the onset of the incident. Around 17,500 companies lost cloud service, including 1,500 premium rate customers who were unfortunate to have their alternate deployment center be one of the other infected sites.

Aftermath Effects - The large majority of CSP A's affected customers had their cloud service restored within 24 hours, but a small proportion were unable to be reconnected for several days,

and for a few, even longer. Some suffered intermittent failures for quite a while. The process of eradication of the worm and the repair of all of the affected systems took several weeks. *A significant portion of the cloud industry was severely impacted in the aftermath of the event, with quite many major CSPs suffering loss of customers as companies experimented with alternatives to third-party cloud service providers. CSP A particularly was badly impacted, losing customers and facing lengthy legal proceedings and lawsuits that took several years to settle.* However, the economics and utility of cloud service to companies ensured that demand for cloud service provision returned to previous and greater levels after some time.

5.2 Estimating Aggregate Cyber-Loss Distributions

The goal of this section is to provide a brief rationale on how cyber-risk estimators (based on an existing real-world instance) from the industry converge upon estimating *aggregate cyber-loss distributions* after a cyber-attack scenario and to consequently show an instance where the aggregate loss distributions derived by these estimators follows a heavy-tailed nature. With respect to the event on which the case study is based, as part of the loss (both first and third party) recovery process, CSP A sought the services of a cyber-insurance firm, whose first step was to send a team of analysts to estimate the loss distribution, $h_L(l) = \mathbb{P}(L \geq l)$. The analysts, after observing relevant parameters after the attack scenario (e.g., commercial losses), arrived at a score, $I_e \in [0, 10]$, indicating the severity of the cyber-attack impact on CSP A. They also derived (using company and loss details) an empirical relationship between I_e and MI, the overall negative impact of the cyber-attack on CSP A's commercial business, and usually a reflection of the combined socio-economic impact of the attack. They arrived at the following deterministic relationship (exact constants sanitized due to company policies):

$$MI = 1.35 + 0.62I_e. \tag{5}$$

Note that this relationship might differ structurally (linear vs. non-linear) for different firms and attack types in practice. In addition, the cyber-risk analysts referred to in the case study, based on their available database of prior cyber-attack information across various companies in their client list, statistically fitted the cyber-attack to generate the following distribution on negative impact M:

$$h_{MI}(m) = C_1 e^{\beta m}, \tag{6}$$

where $\beta = 1.84$, and it is appropriately sanitized for the purpose of this work. Subsequently, from Equation (5), it directly follows that

$$h_{I_e}(i) = C_1 e^{\beta i}.$$

The analysts also estimated the impact intensity (due to cyber-attack percolations through an underlying social communication network, in this case, a service liability network), I_d, at a distance of D hops from the source organizational node in the CSP's liability network graph after the cyber-attack on A to obey the following relation:

$$I_d \geq I_e + 2.9 - 1.25 \log_{10}(D + 10). \tag{7}$$

Clearly the above equation exhibits an attenuation effect of negative impact with distance, and it is somewhat rational and intuitive to believe so, given a single source of attack. Let $A_d(I_e, I_d)$ denote the span (in terms of the number of network (organizational) nodes) that experiences a cyber-attack impact intensity greater than or equal to I_d for a cyber-attack of intensity I_e on the organizational root/source node. Assuming that affected nodes are uniformly distributed across a topological space, we have $A = O(D^2)$. We then have

$$A(I_e, I_d) \geq C_2 10^{1.6(I_e - I_d)} = C_3 e^{1.6 \ln_{10}(I_e - I_d)} = C_4 e^{3.7 I_e} \tag{8}$$

for a fixed I_d. The analysts, using this calculation, consequently were of the opinion that the commercial/economic loss L after the cyber-attack was at least proportional to $A(I_e)$. In addition, the higher the value of I_e, the greater the amount of loss around the source node. *Using these two facets, the cyber-risk analysts mathematically derived the loss distribution to be of the form* $h_L(l) \sim l^{-\alpha}$, *a heavy-tailed distribution*, where $\alpha \in [0.3, 0.76]$.

5.3 Synthetic Evaluation Setup of the Cyber-Insurance Market

Thus far, we have established (courtesy a case study) an important fact: *statistical distributions for aggregated cyber-risk could be heavy tailed.* More specifically, for the case at hand, the cyber-insurer of CSP A after attack analysis came up with a value of $MI = 6.7$ being Bernoulli-Cauchy distributed (based on prior data) with $\mathbb{P}(MI \geq 6.7) = .0063$. Using this specific distributional knowledge from a real-world case (the existence of Bernoulli-Cauchy cyber-risks), we now run synthetic experiments (based on our proposed Bernoulli-Cauchy cyber-risk model) by varying the number of cyber (re-)insurers in a market situation between 5 and 10—the rationale behind which was inferred based on a discussion with CSP A, for the given $MI = 6.7$ value. We assume that the total number of cyber (re-)insurers covers the entire set of organizations affected by a cyber-attack. We allow the maximum number of cyber-risks, N, incurred upon any cyber-insurer to be proportionately distributed in the interval [0,50] according to the *graph density* [Bondy et al. 1976] of the directed network of organization nodes under its coverage, and the maximum coverage liability units on a given absolute scale, k, of each cyber-insurer to uniformly lie in the interval [100, 200]. The rationale behind N stems from the different liability network structures for each organization affected by the cyber-attack, and a higher graph density associates with higher N. The rationale behind choosing k is comparatively more straightforward in the sense that the total loss liability that any cyber-insurer is ready to take up upon itself (on an absolute scale) is based on their market capital that usually lies on an interval in proportion to the liability and is generally not explicitly dependent on the liability network structure. We run multiple iterations of a Bernoulli-Cauchy cyber-risk that hits the source node, and the location parameter, ν, of such a distribution is distributed uniformly in the range $[-10,-20]$ with a scale parameter, σ, fixed to 1. The value of q is uniformly set to lie in the range $[.0005, 0.0063]$ for $MI = 6.7$. The cyber-insurance premiums set by the different market insurers is set in the range of [1,10] loss units distributed proportionately according to the graph density of the directed network of organizational nodes under their coverage, in line with the idea of differentiated pricing based on liability network topologies [Pal et al. 2017]. The degree of risk aversion, α, is set to be uniformly distributed in the interval $[0.03, 0.04]$ based on a discussion with CSP A.

5.4 Evaluation Results

Having established the possibility of heavy-tailed cyber-risks faced by the insurer of a catastrophic cyber-attack source, the main question at hand is: *do individual expected utilities of cyber-risk-pooling insurers at market equilibrium support the formation of a successful cyber-risk re-insurance market?* To this end, we investigate the trends in the expected utility accrued by individual cyber-insurers pooled in a cyber-risk-sharing task, with respect to the number of cyber-risks shared and the number of sharers, i.e., cyber-insurers. We run multiple random instances of cyber-risk attacks on the source node and plot expected utility curves for organizational insurers of arbitrarily chosen liability network nodes without loss of generality, due to the sameness of the plots. *Our goal is to observe expected utility trends, under relaxed assumptions of heterogeneous cyber-risk coverage agents, at market equilibrium and compare the results with our proposed theory derived on the assumption of homogeneous cyber-insurer agents.*

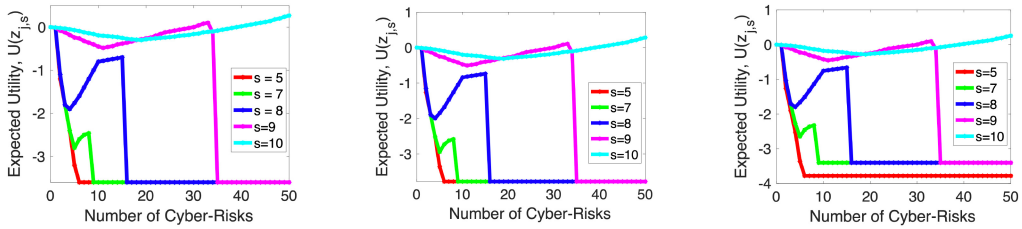

Fig. 1. Cyber-risk sharing for three random instances **[graph density ∈ (0, 0.25)]**.

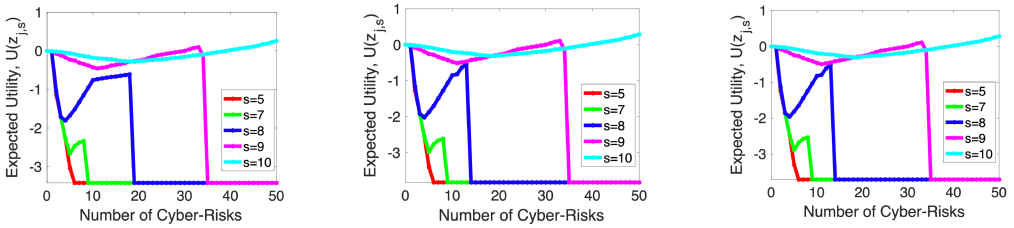

Fig. 2. Cyber-risk sharing for three random instances **[graph density ∈ [0.25, 0.5)]**.

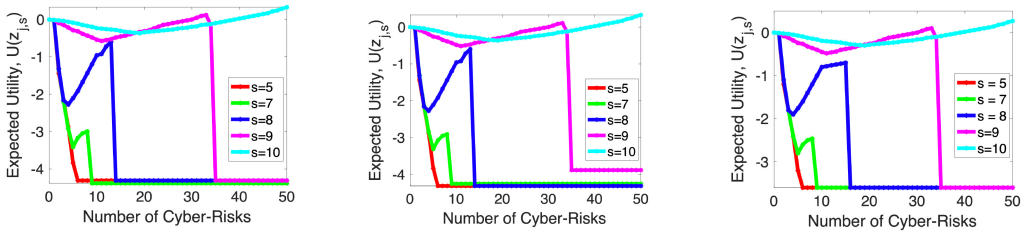

Fig. 3. Cyber-risk sharing for three random instances **[graph density ∈ [0.5, 0.75)]**.

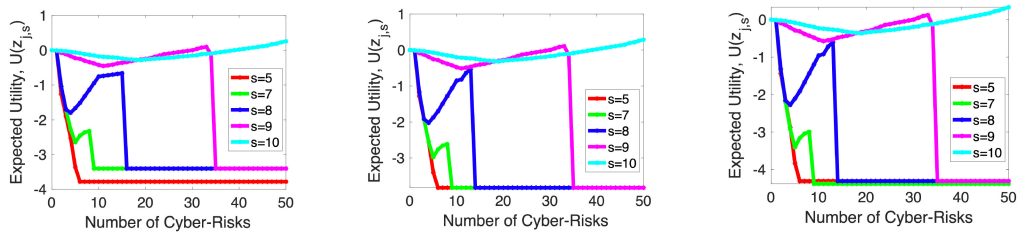

Fig. 4. Cyber-risk sharing for three random instances **[graph density ∈ [0.75, 1)]**.

We observe from the plots in Figures 1 through 4 that for a moderate number of cyber-risks, j, there is no way to increase expected utility compared with staying away from covering cyber-risks altogether (result in accordance with Theorem 5.1). A risk-taker, i.e., cyber-insurer, has the option of not entering the market and must therefore earn a utility premium to be willing to take on cyber-risk (i.e., to offer cyber-insurance). No cyber-insurer will therefore choose (as expected utilities are below zero) to invest in cyber-risks that cannot be pooled. Moreover, if a cyber-insurer believes that no other insurer will pool cyber-risks, it will not take on risks, whether it can pool it or not. Thus, even though the situation with full diversification and cyber-risk sharing $U_{NM,M}$ is preferred over the no cyber-risk situation $U_{0,1}$, at least a threshold number (the s value for which

the expected utility of the cyber-insurer is positive) of cyber-insurers must agree to pool cyber-risk for risk sharing to be worthwhile (result in accordance with Theorems 5.2 and 5.3). In this situation, there may be a coordination problem: even though all the cyber-insurers would like to reach $U_{NM,M}$, they may be stuck in $U_{0,1}$. *Clearly, the limited liability assumption is important to the success of cyber re-insurance markets:* if liability were unlimited, no agent would take on cyber-risk, and diversification of cyber-risk to other cyber-insurers is always inferior (result in accordance with Theorems 5.4 and 5.5). However, we note that the probability for default in the situation with full pooling and cyber-risk diversification is small: approximately around the range of [0.3%, 0.6%] (model based).

6 MATHEMATICAL PROOFS OF RESULTS

In this section, we provide detailed mathematical proofs of our obtained results. Some of the proof constructs are borrowed from Ibragimov et al. [2009].

Proof of Theorem 4.1. The condition C1: $U_{j,s} < U_{0,1}$ for all $j \in \{1, \ldots, N\}$ and all $s \in \{1, \ldots, M\}$ implies $U_{j,1} < U_{0,1}$ for all $j \in \{1, \ldots, N\}$. Thus, it is clearly not optimal for any cyber-insurer who does not participate in the cyber re-insurance market to offer traditional insurance. C1 implies that it is optimal for cyber-insurer m to not offer non-traditional cyber-risk coverage, as any cyber-risk sharing with up to N risks is inferior to not taking on risk. Thus, it is an equilibrium for no one to offer traditional cyber-insurance and is self-evidently unique. □

Proof of Theorem 4.2. If $U_{n,s} > U_{0,1}$ for $n \geq 1$ and $s = 1$, then any cyber-insurer will strictly improve by taking on n cyber-risks. For $U_{n,s} > U_{0,1}$ for some $s > 1$, the proof is a direct consequence of the equilibrium structure of the Stage 2 participation subgame of the RIMG. For example, cyber-insurer 1 strictly improves by pooling n cyber-risks into the cyber re-insurance market, as cyber-insurer $2, \ldots, s^*$ will then choose to participate in the subgame for some $s^* = \text{argmax}_t\{t : U_{n,t} > U_{0,1}\}$, whereas cyber-insurers $s^* + 1, \ldots, M$ will choose not to participate. This leads to a strict improvement for all cyber-insurers $1, \ldots, s^*$. Thus, cyber-insurer 1 will deviate from the assumed strategy of not offering cyber-insurance, and thus this strategy cannot be an equilibrium. □

Proof of Theorem 4.3. Under the conditions provided in the theorem, if cyber-insurer m believes that all other cyber-insurers will participate in the cyber re-insurance market, by choosing N cyber-risks and participating, $U_{NM,M}$ can be achieved. This clearly dominates any alternative strategy of not participating in the market, which will lead to $U_{n,1}$ for $1 \leq n \leq M$, or of participating and offering fewer cyber-risks, which will lead to $U_{n,M}$ for $N(M-1) \leq n \leq NM-1$. All the alternative strategies are thus strictly dominated by the strategy leading to $U_{NM,M}$. □

Proof of Theorem 4.4. The possibility of the existence of multiple (2) RMNEs is shown via an existential argument as follows: for $N = 20$, $M = 5$, $\tilde{X} \in \tilde{S}^q_{\mu,\nu,\sigma}$, along with Cauchy-Bernoulli cyber-risk random variables having parameters $\mu = 1$, $\nu = -9$, $\sigma = 1$, $q = 0.05$, and power utility functions $u(x) = (x + k)^\alpha$, with $k = 100$ and $\alpha = 0.0315$ satisfies (proved below) the following conditions simultaneously:

(1) $U_{j,s} < U_{0,1}$ for all $j \in \{1, \ldots, N\}$ and all $s \in \{1, \ldots, M\}$.
(2) $U_{NM,M} > U_{j,1}$ for all $j \in \{0, \ldots, N\}$.
(3) $U_{NM,M} > U_{j,M}$ for all $j \in \{N(M-1), \ldots, NM-1\}$.

The proof of the satisfiability of these conditions (see below) will lead to two starkly contrasting RMNEs as mentioned in the theorem for $M \geq M_0 = 5$.

We first show that $U_{j,s} < U_{0,1}$ for all $j \in \{1, \ldots, N\}$ and all $s \in \{1, \ldots, M\}$. We do this by studying

$$F(j, y) = \sum_{n=1}^{j} {}^{j}C_n q^n (1-q)^{j-n} W_{j,n,\frac{1}{y}} + (j\mu y + k)^{\alpha},$$

for the parameter values $q = 0.05$, $\mu = 1$, $\nu = -9$, $\alpha = 0.0315$, $\sigma = 1$, $k = 100$. By verifying that $F(j, y) < 0$ for $y \in (0, 1]$, for $j = 1, 2, \ldots, 20$, this implies that $U_{j,s} < U_{0,1}$ for all $s \in \{1, 2, \ldots\}$ and each feasible j. $F(j, y)$ is an increasing function of j for a fixed y, so $F(20, y)$ being strictly negative for $y \in (0, 1]$ is sufficient for condition (1) in the theorem to be satisfied.

To show conditions (2) and (3) in the theorem, let us consider the asymptotics of $U_{tM,M}$, where t is a fixed natural number, as $M \to \infty$. In what follows $\to_P$ denotes convergence in probability and $\to_d$ denotes convergence in distribution. We first mathematically characterize $U_{j,s}$ for the power utility function $u(x) = (x + k)^{\alpha}$, $\alpha \in (0, 1)$, with $k < \infty$ (denoting limited coverage liability) as follows:

$$U_{j,s} = \mathbb{E}\left(\frac{\sum_{i=1}^{j} \tilde{X}_i}{s} + k\right)_+^{\alpha}$$

or

$$U_{j,s} = \mathbb{E}\left(\frac{\mu j + (\nu - \mu)\sum_{i=1}^{j} \varepsilon_i + \sigma \sum_{i=1}^{j} \varepsilon_i Y_i}{M} + k\right)_+^{\alpha},$$

or

$$U_{j,s} = \sum_{n=1}^{j} {}^{j}C_n q^n (1-q)^{j-n} W_{j,n,s} + \left(\frac{j\mu}{s} + k\right)^{\alpha}, \tag{9}$$

where

$$W_{j,n,s} = \frac{1}{\pi} \int_{-\frac{r}{b}}^{\infty} \frac{(bx + r)^{\alpha}}{1 + x^2} dx, \quad b = \frac{n\sigma}{s}, \quad r = k + \frac{(j-n)\mu + n\nu}{s}.$$

The closed-form expression for the integral becomes

$$W_{j,n,s} = \frac{b^{\alpha}(1 + \frac{r}{b})^{\frac{\alpha}{2}}}{\sin(\pi\alpha)} \times \sin\left(\frac{\pi\alpha}{2} + \alpha \tan^{-1}(br)\right).$$

According to Equation (9), we have

$$U_{tM,M} = \mathbb{E}\left(\frac{\sum_{j=1}^{tM} \tilde{X}_i}{M} + k\right)_+^{\alpha}$$

or

$$U_{tM,M} = \mathbb{E}\left(\mu t + (\nu - \mu)\frac{\sum_{i=1}^{tM} \varepsilon_j + \sigma \sum_{j=1}^{tM} \varepsilon_j Y_j}{M} + k\right)_+^{\alpha},$$

where ε_j are i.i.d. nonnegative Bernoulli r.v.s with $P(\varepsilon_j = 0) = 1 - q$, $P(\varepsilon_j = 1) = q$, and $Y_j \in S_{0,1}$ are i.i.d symmetric Cauchy r.v.s with scale parameter $\sigma = 1$ that are independent of ϵ_j's. By the law of large numbers,

$$\frac{\sum_{j=1}^{tM} \varepsilon_j}{M} \to_P t\mathbb{E}\epsilon_1 = tq, \tag{10}$$

as $M \to \infty$. Since the characteristic function of a symmetric Cauchy r.v. $X \in S_{0,\sigma}$ is given by

$$\mathbb{E}\exp(iyX) = \exp(-\sigma|y), \tag{11}$$

we obtain that the characteristic function $f(y) = \mathbb{E}\exp(iyW_M)$ of $W_M = \frac{\sum_{j=1}^{tM} \varepsilon_j Y_j}{M}$ satisfies

$$f(y) = \left[\mathbb{E}\exp\left(\frac{iy\epsilon_1 Y_1}{M}\right)\right]^{tM} = \left[1 + q\left(\exp\left(-\frac{|y|}{M}\right) - 1\right)\right]^{tM} \to \exp(-qt|y)\ y \in \mathrm{R}, \tag{12}$$

as $M \to \infty$. Because, according to Equation (11), $\exp(-tq|y)$ is the characteristic function of the r.v., tqY_1, we conclude from Equation (12) that

$$W_M \to_d tqY_1, \tag{13}$$

as $M \to \infty$. Relations (13) and (10) imply that, as $M \to \infty$,

$$V_M = \left(\mu t + (\nu - \mu)\frac{\sum_{i=1}^{tM} \varepsilon_j + \sigma \sum_{j=1}^{tM} \varepsilon_j Y_j}{M} + k\right)_+^{\alpha} \to_d (\mu t + (\nu - \mu)tq + \sigma tqY_1 + k)_+^{\alpha}. \tag{14}$$

Because, as it is not difficult to see, the sequence of r.v.s V_Ms is uniformly integrable, from Equation (14) we get (from a result in Ash et al. [2000]) that

$$U_{tM,M} = \mathbb{E}\left(\mu t + (\nu - \mu)\frac{\sum_{i=1}^{tM} \varepsilon_j + \sigma \sum_{j=1}^{tM} \varepsilon_j Y_j}{M} + k\right)_+^{\alpha}$$

converges to

$$G(t) = \mathbb{E}(\mu t + (\nu - \mu)tq + \sigma tqY_1 + k)_+^{\alpha},$$

or

$$G(t) = \left(t^2q^2\sigma^2 + (k + t(\mu - q\mu + q\nu))^2\right)^{\frac{\alpha}{2}} \times \csc(\pi\alpha)\sin\left(\frac{1}{2}\alpha\left(\pi + 2\tan^{-1}\left(\frac{k + t(\mu - q\mu + q\nu)}{qt\sigma}\right)\right)\right),$$

as $M \to \infty$. In particular,

$$\lim_{M\to\infty} U_{nM,M} = G(n),\ n = 1, \ldots, N.$$

It is straightforward to see that $G(N) > U_{0,1}$.

Let us now show that $U_{N(M-1),M} < U_{N(M-1)+1,M} <, \ldots, U_{NM-1,M}$. Let $0 \le m \le M - 2$. Note that, for $Y \in S_{0,\sigma}$ and $\alpha \in (0, 1)$, $z \in \mathrm{R}$, the expectation $\mathbb{E}(Y + z)_+^{\alpha-1}$ is finite and well defined because the integral $\int_{-z}^{\infty} \frac{dx}{\pi\sigma(x+z)^{1-\alpha}(1+x^2)}$ converges. This, by induction and conditioning arguments, implies that the expectation $\mathbb{E}(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k)_+^{\alpha-1}$ is finite and positive, i.e.,

$$0 < \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha-1} < \infty. \tag{15}$$

Using a Taylor expansion, it is not difficult to check that the following inequality holds for all $x, z \in \mathrm{R}$, and for all $\alpha \in (0, 1)$:

$$(x + z)_+^{\alpha} \ge z_+^{\alpha} + \alpha x z_+^{\alpha-1}, \tag{16}$$

with strict inequality for $x, z > 0$. Let $X_1' = \mu + (\nu - \mu)\varepsilon_1' + \sigma\varepsilon_1' Y_1'$ denote an r.v. with Bernoulli-Cauchy distribution, where ε_1' is a non-negative Bernoulli r.v. with $P(\varepsilon_i = 0) = 1 - q, P(\varepsilon_i = 1) = q$, and $Y_1' \in S_{0,1}$ is a symmetric Cauchy r.v. with scale parameter $\sigma = 1$ independent of ε_1'. Suppose

further that ε_1' and Y_1' are independent of the r.v.s $X_j, j = 1, \ldots, N(M-1)+1$. Using Equations (15) and (16) we get that, for all $s > 0$,

$$\mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j + X_1'}{M} + k\right)_+^{\alpha} I(|Y_1'| < sM) > \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha} P(|Y_1'| < sM) + \left(\frac{\alpha}{M}\mathbb{E}[X_1' I(|Y_1'| < sM)]\mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha-1}, \tag{17}$$

where $I(\cdot)$ is the indicator function. We further have, by the definition of X_1', ε_1', and Y_1' and using the symmetry of the distribution of Y_1',

$$\mathbb{E}[X_1' I(|Y_1'| < sM)] = (\mu + (\nu - \mu)q)P(|Y_1'| < sM) + \sigma q \mathbb{E}[Y_1' I(|Y_1'| < sM)] \tag{18}$$

or

$$\mathbb{E}[X_1' I(|Y_1'| < sM)] = (\mu + (\nu - \mu)q)P(|Y_1'| < sM),\ \forall s > 0.$$

The inequalities in Equations (17) and (18) imply that, for all $s > 0$,

$$\mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j + X_1'}{M} + k\right)_+^{\alpha} I(|Y_1'| < sM) > \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha} P(|Y_1'| < sM) + \frac{\alpha(\mu + (\nu - \mu)q)}{M} P(|Y_1'| < sM). \tag{19}$$

Letting $s \to \infty$, we obtain

$$U_{N(M-1)+m+1} = \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j + X_1'}{M} + k\right)_+^{\alpha},$$

or

$$U_{N(M-1)+m+1} = \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha} + \frac{\alpha(\mu + (\nu - \mu)q}{M}. \tag{20}$$

It is easy to check that, for the values of the parameters chosen in the theorem, $\mu + (\nu - \mu)q = 0.5 > 0$. From Equation (20) we thus conclude that

$$U_{N(M-1)+m+1} = \mathbb{E}\left(\frac{\sum_{j=1}^{N(M-1)+m} X_j}{M} + k\right)_+^{\alpha} = U_{N(M-1)+m}$$

for all $0 \le m \le M-2$. Thus, the conditions for a risk-sharing diversification equilibrium are satisfied for all $M \ge M_0$, hence proving the theorem. □

Proof of Theorem 4.5. Assume that there exists a two-equilibrium state to an instance, (u, F, ∞, N, M), of the RIMG for arbitrarily large M, for a strictly concave, twice continuously differentiable utility function u and distribution F satisfying

$$\int_{-\infty}^{\infty} x^2 dF = C < \infty. \tag{21}$$

For notational convenience, we assume that F is differentiable, so that with pdf $\phi = \frac{dF}{dx}$ is well defined. However, the whole proof goes through step by step without this restriction. Without

loss of generality, we can assume $u(0) = 0, u'(0) = 1$. For a genuine two-equilibria state to exist, it must be the case that for arbitrarily large M, we have

$$U_{NM,M} > 0, \tag{22}$$

and

$$U_{1,M} < 0. \tag{23}$$

However, a necessary condition for Equation (22) to hold for arbitrarily large M is that $\mathbb{E}[X] \leq 0$, as seen by the following argument: assume that $\mu = \mathbb{E}[X] > 0$. We define

$$U(\varepsilon) = \int_{-\infty}^{\infty} u(\epsilon x)\phi(x)dx.$$

We decompose

$$u(x) = x - t(x) = x - x^2 z(x),$$

where $t(0) = t'(0) = 0, t'' > 0, t(x) < x, z(x)$ is continuous, and both t and z are non-negative. We then have

$$U(\epsilon) = \epsilon\mu - \left(\int_{-\infty}^{-\frac{1}{\varepsilon}} t(\varepsilon x)\phi(x)dx + \int_{-\frac{1}{\varepsilon}}^{\frac{1}{\varepsilon}} (\varepsilon x)^2 z(\varepsilon x)\phi(x)dx + \int_{\frac{1}{\varepsilon}}^{\infty} t(\varepsilon x)\phi(x)dx \right). \tag{24}$$

The $\int_{\frac{1}{\varepsilon}}^{\infty}$ term is clearly $o(\varepsilon)$, as

$$\int_{\frac{1}{\varepsilon}}^{\infty} t(\varepsilon x)\phi(x)dx \leq \int_{\frac{1}{\varepsilon}}^{\infty} \varepsilon x\phi(x)dx \leq \varepsilon \int_{\frac{1}{\varepsilon}}^{\infty} x\phi(x)dx \leq C_2\varepsilon^2.$$

Furthermore, as $z(x)$ is continuous, it is bounded on $[-1, 1]$, so Holder's inequality can be used to bound the $\int_{\frac{1}{\varepsilon}}^{\frac{1}{\varepsilon}}$ term by

$$\int_{\frac{1}{\varepsilon}}^{\frac{1}{\varepsilon}} (\varepsilon x)^2 x(\varepsilon x)\phi(x)dx \leq \varepsilon^2 \max_{-1 \leq y \leq 1} |z(x)| \times C = C_3\varepsilon^2,$$

so the second term is also of $o(\varepsilon)$. Finally, the $\int_{-\infty}^{-\frac{1}{\varepsilon}}$ term is also $o(\varepsilon)$, as

$$\int_{-\infty} -\frac{1}{\varepsilon} t(\varepsilon x)\phi(x)dx = \int_{infty}^{-\frac{1}{\varepsilon}} \frac{t(\varepsilon x)}{t(x)} t(x)\phi(x)dx \leq \varepsilon t'(-1) \times \int_{-\infty}^{-\frac{1}{\varepsilon}} t(x)\phi(x)dx = o(\varepsilon),$$

where we use Holder's inequality to move the $\frac{t(\epsilon x)}{t(x)}$ outside of the integral, and the inequality

$$\frac{t(\varepsilon x)}{t(x)} \leq \varepsilon t'(-1),$$

which must hold for $x \geq \frac{1}{\varepsilon}$, as t is convex. Finally,

$$\int_{-\infty}^{-\frac{1}{\varepsilon}} t(x)\phi(x)dx = o(1),$$

as the integral $\mathbb{E}u(X)$ could otherwise not exist. This altogether implies that $U(\varepsilon) = \varepsilon\mu - o(\varepsilon)$, which is strictly positive for small enough ε. Therefore, if $\mathbb{E}(X) > 0$, then $U_{1,M}$ will be strictly positive for large enough M, and no genuine two-equilibrium state can therefore exist. However, if $\mathbb{E}(X) \leq 0$, then a two-equilibrium state cannot exist, as Jensen's inequality implies that $U_{NM,M}$ is strictly negative for arbitrary $M > 0$ and $N > 0$ and thus $U_{NM,M} < U_{0,1}$. This proves the theorem for clause (i). Now assume that there exists a two-equilibrium state to the game (u, ϕ, k, N, M) for

arbitrarily large M, for a strictly concave, twice continuously differentiable utility function u and distribution ϕ, satisfying

$$\int_{-\infty}^{\infty} x^2 \phi(x) dx = C < \infty. \tag{25}$$

Without loss of generality, we can assume $u(0) = 0, u'(0) = 1$. If $\mathbb{E}(X) > 0$, then the same argument as in the proof of the first part of Theorem 4.5 rules out the two-equilibria state for $M \geq M_0$, as the limited liability increases $U_{1,M}$ compared to the unlimited liability case. Thus, for M large enough, $U_{1,M}$ must be strictly positive and a two-equilibrium state cannot exist. If $\mathbb{E}[X] = \mu < 0$, then we can use the law of large numbers to show that as M becomes large, $X_{NM} = (NM)^{-1} \sum_{i=1}^{NM} X_i$ converges in distribution to μ. Thus, $\lim_{M\to\infty} \mathbb{E}((X_{NM} + kNM)_+ - kNM) = \mu$, so for some large enough M_0, $\lim_{M\to\infty} \mathbb{E}((X_{NM} + kNM)_+ - kNM) < 0$ for all $M \geq M_0$. Jensen's inequality therefore again implies that $U_{NM,M}$ is strictly negative for $M \geq M_0$ and thus $U_{NM,M} < U_{0,1}$, so there cannot be a two-equilibria state, thus proving the theorem regarding part (ii) of the theorem clause. We now prove that $U_{NM,M} < U_{0,1} = 0$ for large M. We define

$$\gamma = \min_{x\in[-\frac{k}{2},\frac{k}{2}]} u''(x).$$

As u is strictly concave and twice continuously differentiable, $\gamma > 0$. We define

$$\tilde{u}_x = x - \frac{\gamma}{2} x^2 I_{[-\frac{k}{2},\frac{k}{2}]},$$

where I_A is the indicator function on the set A, implying that $u(x) \leq \tilde{u}(x)$ for all $x \in [-k, \infty)$. Similar to $U_{j,s}$, we define $\tilde{U}_{j,s}$, the "utility" of sharing j cyber-risks equally among s cyber-insurers, for insurers with "utility" functions $\tilde{u}$. Clearly $U_{j,s} \leq \tilde{U}_{j,s}$, so if $\tilde{U}_{NM,M} < 0$ for large M, then $U_{j,s} < 0$ for large M and there cannot be a two-equilibria state for an M_0 such that $M \geq M_0$. We next define $Y_1 = \sum_{i=1}^{N} X_i$ and study uniform portfolios of i.i.d. cyber-risks $Y_1, \ldots, Y_M$ by defining $\bar{Y}_M = \frac{\sum_{m=1}^{M} Y_M}{M}$. As $\mathbb{E}(\bar{Y}_M) = 0$, the condition $\tilde{U}_{NM,M} < 0$ for large M can be written as

$$\tilde{U}_{NM,M} = \mathbb{E}(\bar{Y}_M I_{[-k,\infty)}) - \frac{\gamma}{2} \mathbb{E}(\bar{Y}_M^2 I_{[-\frac{k}{2},\frac{k}{2}]}) < 0. \tag{26}$$

We begin by bounding $\mathbb{E}(\bar{Y}_M^2 I_{[-\frac{k}{2},\frac{k}{2}]})$ from below. From the central limit theorem, we know that $Z_M = \sqrt{MY_M}$ converges in distribution to $Z \sim \mathcal{N}(0, \sigma^2)$, so $\mathbb{E}(Z_M^2 I_{[-\frac{k}{2},\frac{k}{2}]}) \to C > 0$, as M grows. Now since

$$M\mathbb{E}(\bar{Y}_M^2 I_{[-\frac{k}{2},\frac{k}{2}]}) \geq M\mathbb{E}(\bar{Y}_M^2 I_{[-\frac{k}{2\sqrt{M}},\frac{k}{2\sqrt{M}}]}) = \mathbb{E}(Z_M^2 I_{[-\frac{k}{2},\frac{k}{2}]}]),$$

we can conclude that for large M,

$$\frac{\gamma}{2} \mathbb{E}(\bar{Y}_M^2 I_{[-\frac{k}{2},\frac{k}{2}]}) \geq \frac{C'}{M}, \; C' > 0. \tag{27}$$

We next bound $\mathbb{E}(\bar{Y}_M I_{[-k,\infty)})$ from above. As $\mathbb{E}(\bar{Y}_M) = 0$, we have $\mathbb{E}(\bar{Y}_M I_{[-k,\infty)}) = -\mathbb{E}(\bar{Y}_M I_{[-\infty,-k)})$. From the Cauchy-Schwarz inequality, we know that

$$-\mathbb{E}(\bar{Y}_M I_{[-\infty,-k)}) \leq \mathbb{E}\left(\bar{Y}_M^2\right)^{\frac{1}{2}} \mathbb{E}(I_{(-\infty,-k)})^{\frac{1}{2}}$$

(as $I^2_{(-\infty,-k)} = I_{(-\infty,-k)}$). Of course, $\mathbb{E}(\bar{Y}_M^2) = \frac{\sigma^2}{M}$. Moreover, Rosenthal's inequality [Rosenthal 1973] implies that $\mathbb{E}(Y_M^{2+\epsilon}) \leq \frac{C''}{M^{1+\frac{\epsilon}{2}}}$, and by Markov's inequality [Ross 2014], we therefore know that

$$\mathbb{E}(I_{(-\infty,-k)}) = P(x < -k) \leq \frac{\mathbb{E}(Y_M^{2+\epsilon})}{k^{2+\epsilon}} \leq \frac{C'' k^{-(2+\epsilon)}}{M^{1+\frac{\epsilon}{2}}}.$$

Overall, this implies that

$$\mathbb{E}(\bar{Y}_M I_{(-k,\infty)}) \leq \sqrt{\frac{\sigma^2}{M}}\sqrt{\frac{C''k^{-2+\epsilon}}{M^{1+\frac{\epsilon}{2}}}} = \frac{C'''}{M^{1+\frac{\epsilon}{4}}}.$$

The bounds in Equation (27) are therefore

$$\tilde{U}_{NM,M} \leq \frac{C'''}{M^{1+\frac{\epsilon}{4}}} - \frac{C'}{M}, C' > 0,$$

which is strictly negative for large M. Thus, as $U_{NM,M} \leq \tilde{U}_{NM.M}$, we know that $U_{NM,M} \leq U_{0,1} = 0$ for large M. Therefore, there cannot be any two-equilibria state that can arise in this case either, hence proving the theorem for case (iii). □

Proof of Theorem 4.7. It is evident from the scenario of non-curtailed cyber-risk with finite support (see Theorem 3 in Pal et al. [2020b]) that for any $t > 0$, we can choose an a such that

$$F_{Y_1(a)}(-t) < F_{\bar{Y}_n(a)}(-t).$$

We have

$$\int_{-\infty}^{-t} -t dF_{Y_1(a)}(x) + \int_{-t}^{\infty} x dF_{Y_1(a)}(x) = \int_{-\infty}^{-t} -t dF_{Y_n(a)}(x) + \int_{-t}^{\infty} x dF_{Y_n(a)}(x) - s, \quad s > 0.$$

$s > 0$ follows from Rothschild and Stiglitz [1970] and that $Y_1(a)$ is a mean preserving spread of $\bar{Y}_n(a)$. Specifically, the integral takes the form $\int QdF$, where $Q(x) = (x+t)_+ - t$ is convex and therefore $-Q$ is concave. For a specific a, we can clearly choose an admissible utility function u, such that

$$\left|\int_{-t}^{\infty} u(x) dF_{Y_1(a)}(x) - \int_{-t}^{\infty} x dF_{Y_1(a)}(x)\right| \leq \frac{s}{6}, \quad \left|\int_{-t}^{\infty} u(x) dF_{Y_n(a)}(x) - \int_{-t}^{\infty} x dF_{Y_n(a)}(x)\right| \leq \frac{s}{6},$$

and, for t large enough,

$$\left|\int_{-\infty}^{-t} v(x) dF_{Y_1(a)}(x) - \int_{-\infty}^{-t} -t dF_{Y_1(a)}(x)\right| \leq \frac{s}{6}, \quad \left|\int_{-\infty}^{-t} v(x) dF_{\bar{Y}_n(a)}(x) - \int_{-\infty}^{-t} -t dF_{\bar{Y}_n(a)}(x)\right| \leq \frac{s}{6}.$$

We therefore have

$$\begin{aligned}
&\int_{\infty}^{\infty} v(x)\left[dF_{Y_1(a)}(x) - dF_{\bar{Y}_n(a)}(x)\right] = \int_{-\infty}^{\infty} (Q(x) + v(x) - Q(x))\left[dF_{Y_1(a)}(x) - dF_{\bar{Y}_n(a)}(x)\right] \\
&= s + \int_{-\infty}^{-t} (v(x) - (-t))\left[dF_{Y_1(a)}(x) - dF_{\bar{Y}_{n(a)}}(x)\right] + \int_{-t}^{\infty} (u(x) - x)\left[dF_{Y_1(a)}(x) - dF_{\bar{Y}_n(a)}(x)\right] \\
&\geq s - \left|\int_{-\infty}^{-t} (v(x) - (-t)) dF_{Y_1(a)}(x)\right| - \left|\int_{-\infty}^{-t} (v(x) - (-t)) dF_{\bar{Y}_n(a)}(x)\right| \\
&\quad - \left|\int_{-t}^{\infty} (u(x) - x) dF_{Y_1(a)}(x)\right| - \left|\int_{-t}^{\infty} (u(x) - x) dF_{\bar{Y}_n(a)}(x)\right| \\
&\geq s - 4\frac{s}{6} = \frac{s}{3}.
\end{aligned}$$

Altogether, this implies that

$$Ev\left(Y_1(a)\right) \geq Ev(\bar{Y}_n(a)) + \frac{s}{3},$$

and as $s > 0$, we are through. □

7 RELATED WORK

In this section, we cite other works most related to ours in this article. However, we would like to emphasize upfront that a formal analysis on the *economic incentive compatibility of cyber-reinsurance services under heavy-tailed (i.i.d. in our work) cyber-risk distribution type* is absent in (IoT) literature for cyber-insurance or network risk management settings, and our efforts here in this direction are completely new to the best of our knowledge. In two very recent efforts [Pal et al. 2020c, 2020a], we proposed an orthogonal formal analysis on the *statistical incentive compatibility of providing cyber-reinsurance services,* the first of its kind, in IoT societies under complete general cyber-risk distribution types be it i.i.d./non-i.i.d. heavy tailed or otherwise. The analysis in these papers, unlike here, is oblivious to re-insurer rationality while making service feasibility decisions on risk spreading. However, they account for the *complete general nature* of statistical distributions in the decision-making process.

7.1 Success of Cyber-Insurance Markets

In this work we investigated the feasibility of a pool of cyber-insurers getting together to cover aggregate cyber-risks of a heavy-tailed statistical nature. However, a pre-cursor to having aggregate risk-covering cyber re-insurance markets is working successful primary cyber-insurance markets in the first place. To this end, recent research works on cyber-insurance [Hoffman 2007; Lelarge and Bolot 2009b; Shetty et al. 2010] have mathematically shown the existence of economically inefficient insurance markets. Intuitively, an efficient market is one where all stakeholders (market elements) mutually satisfy their interests. These works state that cyber-insurance satisfies every stakeholder apart from the regulatory agency (e.g., government), and sometimes the cyber-insurer itself. The regulatory agency is unsatisfied as overall network robustness is sub-optimal due to network users not optimally investing in self-defense mechanisms, whereas a cyber-insurer is unsatisfied due to it potentially making zero expected profit at times. In Pal and Golubchik [2010], the authors proposed a Coasian bargaining approach among cyber-insured network entities to achieve an efficient insurance market; however, costless bargaining under which the Coase theorem holds is idealistic in nature and might not be feasible to implement in practice. Lelarge and Bolot [2009b] recommended the use of fines and rebates on cyber-insurance contracts to make each user invest optimally in self-defense and make the network optimally robust. However, their work neither mathematically proves the effectiveness of premiums and rebates in making network users invest optimally nor guarantees the strict positiveness of insurer profits at all times. In recent works [Pal et al. 2011, 2014, 2018; Khalili et al. 2018], the authors overcome the drawbacks of the mentioned existing works and propose ways to form provably efficient monopolistic cyber-insurance markets by satisfying market stakeholders, including a risk-averse cyber-insurer, in environments of interdependent risk. Naghizadeh and Liu [2014], Pal et al. [2011], and Naghizadeh and Liu [2016] further state the importance of compulsory insurance for optimizing social welfare for primary cyber-insurance markets. In addition, and more importantly, recent major successful cyber-attacks on large commercial organizations have significantly increased board-level concerns to maintain business reputation amongst clients and subsequently accelerated the adoption of cyber-insurance products in the industry.

Drawbacks - These works are completely orthogonal in their goals, when compared to our effort (i.e., economically investigating primary cyber-insurance market success vs. statistically investigating the effect of heavy-tailed cyber-risks on the existence of aggregated insurance coverage markets), and do not investigate the effect of aggregated cyber-risk on the inclination of a cyber-insurer to participate in (insurance/re-insurance) coverage markets for a service-networked setting—a prime determinant for the expansion of the insurance industry for the modern cyber-age.

7.2 On the Heavy-Tailed and Dependent Nature of Cyber-Risk

There are quite a few *instances in the practical real world* where cyber-risks have shown heavy-tailed impact. Maillart and Sornette [2010] analyzed a *Datalossdb* 2017 dataset consisting of 956 personal identity loss incidents that occurred in the United States between the years 2000 and 2008. They found that the personal identity losses per incident, denoted by X, can be modeled by a heavy-tail distribution $P(X > n) \sim n^{\alpha}$, where $\alpha = 0.7 +/- 0.1$, and more importantly, this result holds for a variety of organizations: business, education, government, or a medical institution. Because the probability density function of the identity losses per incident is static, the situation of identity loss is stable from the point of view of the breach size. Edwards et al. [2016] analyzed a Privacy Rights Clearinghouse database of 2017 consisting of 2,253 breach incidents that spanned over a decade from 2005 to 2015. These breach incidents include two categories: negligent breaches (i.e., incidents caused by lost, discarded, stolen devices or other reasons) and malicious breaching (i.e., incidents caused by hacking, insider, and other reasons). They showed that the breach size can be modeled either by log-normal or log-skew-normal distribution that are not heavy-tailed distributions in a mathematically precise sense, but have long tails, or by Pareto distributions that are heavy tailed. Wheatley et al. [2016] merged and analyzed cyber-breach incidents from the Datalossdb and the Privacy Rights Clearinghouse database spanning over a decade (2000 to 2015). They used the **Extreme Value Theory (EVT)** [Embrechts et al. 2013] to study the maximum breach size and further modeled the large breach sizes by a doubly truncated heavy-tailed Pareto distribution. There are also studies establishing the dependence among cyber risks. Notable among them are Herath and Herath [2011], Pal et al. [2019, 2020c], Mukhopadhyay et al. [2013], Böhme and Kataria [2006], Xu et al. [2017, 2018], Xu and Hua [2019], and Peng et al. [2018].

Drawbacks - Existing empirical research in cyber-security has been successful in elucidating the heavy-tailed and tail-dependent nature of cyber-risk; however, *it is yet to propose formally proven directions to allow a profit-minded cyber-risk manager to judge whether a collection of such risks is suitable to aggregate, under various degrees of heavy-tailedness.* This decision-making problem will increasingly arise in the IoT age where major cyber-risks affecting smart societies will give rise to systemic effects that cyber-risk managers have to deal with. It is a common perception from empirical studies and traditional insurance literature [Samuelson 1967] that i.i.d. non-catastrophic cyber-risks, even though heavy tailed, are suitable for aggregation. *In this article, as a first, we showed quite the contrast for i.i.d. catastrophic heavy-tailed risks.*

7.3 Analyzing Network Effects of Cascading Events upon Cyber-Insurers

The typical cyber re-insurance scenario will arise in a service network setting where insurance companies covering individual organizations are willing to diversify their client cyber-risks among multiple other cyber-insurers. The liability network effects when such a network of cyber-insurers is hit via coverage demands after a catastrophic cyber-attack are a function of (1) the probability of an organization being hit by cyber-attacks and (2) the effect aggregate cyber-risks can have on potential risk-covering organizations, e.g., re-insurers. Nearly all research (to the best of our knowledge) in the information security community has only works in (1), i.e., mathematical/simulation models for estimating the probability of a consumer (e.g., organization, end-user) getting infected by cyber-attack vectors (e.g., APTs) (see survey in Böhme et al. [2010]) [Böhme et al. 2017]. If we denote such a probability distribution for each consumer i by F_i, we emphasize that evaluating F_i actually involves mathematically capturing the spread of the infection (attack) vector (e.g., a virus, bot), and is not the focus of this article. The interested reader is referred specifically to Lelarge and Bolot [2009b], Lorenz et al. [2009], and Ganesh et al. [2005] to get insights on statistical mean field models to mathematically and robustly evaluate $F()$. Regarding the costs faced by the consumers

on getting cyber-attacked, the readers are referred to Riek and Böhme [2018]. In terms of the process of the physical spread of attacks in networks, a related literature has directly originated from the study of cascades. Various models have been developed in the computer science and network science literatures, including the widely-used threshold models [Granovetter 1978] and percolation models [La 2016, 2018a, 2018b; Watts 2002; Molloy and Reed 1998, 1995; Newman et al. 2001; Chung and Lu 2002]. A few works have applied these ideas to various economic settings, including Durlauf [1993] and [Bak et al. 1993] in the context of economic fluctuations; [Morris 2000] in the context of contagion of different types of strategies in coordination games; and more recently, [Gai and Kapadia 2010] and [Blume et al. 2011] in the context of spread of an epidemic-like financial contagion, where the seminal papers of [Allen and Gale 2000] and [Freixas et al. 2000] developed some of the first formal models of contagion over financial networks.

The only work to the best of our knowledge that has considered the effect of the insurance network and aggregate cyber-risks on cyber-insurance companies is the contribution by Pal et al. [2020d]. More specifically, the authors conduct a robust network analysis and surprisingly find that the specific underlying service network and a broad family of statistical cyber-risk distributions (including heavy-tailed distributions) did not have "much" of an effect (the network in general, though, plays a role) on the increase in the amount of "to be covered expected loss" in the aftermath of a catastrophic cyber-attack. What primarily mattered was (1) the revenue base of the service-providing organization and (2) the extent of reliance of an organization upon others. This result is beneficial and impactful in the sense that cyber-insurers mostly do not have complete knowledge of the underlying topology that ties service organizations together, and it is equally difficult to get a precise estimate of cyber-risk distributions.

Drawbacks - What we are concerned about with respect to contributions made in the article is part (2), i.e., the investigation (given that part (1) as above-mentioned is taken care of) of how aggregate cyber-risk arising from multiple network nodes, i.e., cyber-insurers covering individual risks of organizations in the event of a cyber-attack, affects the incentive compatibility mindset of the aggregate risk-covering node. To the best of our knowledge, none of the existing works have tackled this issue statistically and economically (as we have done), and this is precisely what is important when considering whether a market for cyber re-insurance will thrive in the modern cyber-age. As mentioned above, Pal et al. [2020d] give a statistical and network-oblivious perspective to the problem at hand, *but they do not comment on the incentive compatibility of covering aggregate cyber-risk.* They only claim that the *specifics* underlying connected service network topology, in addition to the cyber-risk distributions, do not contribute relatively significantly to the increase in the expected cyber-losses in the aftermath of a catastrophic cyber-attack.

8 DISCUSSION

In this section, we first provide a stance about the possibility of commercial risk management organizations encountering heavy-tailed cyber-losses in the IoT age, as viewed by some corporate leaders in the current cyber-insurance market space. Second, we briefly describe the state of current cyber re-insurance markets. Finally, we discuss the role of regulation in preventing societal unwanted situations where profit-minded cyber re-insurers shy away, i.e., do not feel incentive compatible, from providing re-insurance services for the benefit of society even if they can, in this hour of need.

8.1 Industry Expert Views on Cyber-Insurance in the IoT Age

According to Pascal Millaire, a cyber-insurance expert and the chief executive officer of *CyberCube*, the insurance industry has a long way to go to meet cyber coverage business interruption needs. He noted, *"I think today we're only scratching the surface when it comes to the need for*

business interruption cover in the cyber space," agreeing that the IoT will push the issue to a crisis point. The IoT age will unlock over USD 10 trillion in new economic activity globally with a projected 50 billion IoT devices to be in operation by 2025. While many of these devices will work well and benefit society, many will be vulnerable even to naive cyber-attacks. To make things worse, due to common software platforms and applications running on many of these devices, the attack-impact effects will be correlated in time and space. Consequently, when these fail, it will not be unreasonable to expect significant aggregate cyber-loss impacts with statistically very heavy tails, simply because with increasing inter-connectivity, the cost of a failed connection goes higher. According to data released by UK regulatory authorities in 2019, enterprise-related cyber-losses are ranked higher in the country when compared to flood, Japan earthquakes, and most other catastrophes other than California earthquakes and Atlantic hurricane-related losses. Michael Tannenbaum, executive vice president for *Financial Lines* for *Chubb,* believes that cyber-insurance coverage can address and mitigate such impacts via adding services such as call centers, threat notifications, surge protection, IT restoration, and legal services that help clients counter a threat quickly and mitigate long-term damages. At the same tine, cyber-insurers need to be able to facilitate hundreds of claims at the same time when business interruption-related cyber-attacks strike.

8.2 On Current Cyber (Re)-Insurance Markets

Re-insurance is still a nascent market with emerging risks that are yet to be priced or modeled adequately (courtesy SCOR Report, 2017 [Coburn et al. 2018]). Because cyber re-insurers concentrate the risks ceded to them by insurance companies, the aggregation of cyber exposure coming from cyber-specific or standard products is exacerbated. Developing modeling capabilities to get a grip on clash of risks and cyber catastrophes is a condition for the cyber re-insurance market to grow. The current global reinsurance market is estimated to be worth approximately USD 525M. Most cyber re-insurers have only just entered the cyber insurance market. Because of the modeling and pricing issues mentioned above, it remains mostly a proportional market with approximately 95% of cyber re-insurance premiums being written by re-insurers on a quota share basis. More importantly, unlike other lines of insurance, cyber has only a short history of experience, and actuarial analysis is made more complicated by rapid changes in the threat and loss patterns from year to year. Instead, cyber-insurance companies have sold policies that represent relatively limited exposure to themselves, chiefly through constraining the level of limit that they provide. An estimated half of all cyber insurance policies sold are for limits of less than $1 million, i.e., the total amount that insurers are prepared to pay out from any cyber event is capped at $1 million. Limits of over $10 million are rare (less than 10% of policies written), and for a company to obtain cyber insurance coverage of $100 million (quite a possible scenario for cyber re-insurance) or more requires the construction of complex "towers" of coverage involving many different insurance companies, each taking a small slice. Limits are increasing over time as cyber-insurers gain confidence, but the protection being offered is not what is being requested by the market. Things could get worse liability-wise in the future densely device-populated era of the IoT with (cloud-connected) IoT environments like smart-homes, smart (water) grid, smart agriculture, etc., being increasingly reliant upon cyber-insurance companies in situations of service failures. Add to this the extreme risk-averse mindset of cyber-insurers towards cyber-catastrophes and the thought of future cyber-attack scenarios such as a worse version of the 2017 *NotPetya* malware, instead of causing multi-million-dollar losses to several dozen corporations, hitting thousands of companies, triggering full-limit claims from a sizeable proportion of the insurer's portfolio. In such cases the cyber-insurer could have run a cyber insurance business for a decade profitably, achieving low loss ratios, and then have a single year in which all the reserves it has built up or more are wiped

out. The frequency and the severity of these multiple-claim catastrophes determine the long-term profitability and viability of cyber risk as a line of insurance business, and through our analysis in this article, we show that such multiple-claim catastrophes may indeed be not incentive compatible for a flourishing cyber re-insurance industry, unless government regulation plays an important role.

8.3 The Role of a Regulator

It is well known [Bruggeman et al. 2010] that the private catastrophe insurance markets for earthquakes, wind damage, floods, terrorism, and the current COVID-19 outbreak have failed (in our work the analogous state of providing no cyber re-insurance coverage on a cyber-catastrophe), one after the other, over the last 25 years. As a result, in both the United States and Europe, governments have been forced to intervene and design risk management plans. The plans were often created under time pressure and they differ substantially in their details; see OECD measures in Bruggeman et al. [2010] for descriptions of both the European and U.S. plans. However, *it is intriguing to find that they actually share a fundamental design feature, namely that each government plan has, in effect, created a mechanism through which, similar to our results, a coordinated equilibrium is established, where providing re-insurance is incentive compatible. We envision and recommend such a similar intervention to boost the success of cyber (re-)insurance markets.* For example, an act could be passed for cyber-insurance firms to offer cyber-catastrophe coverage as a rider to their standard primary coverage. A *quid pro quo* arrangement could be worked out where the federal government can provide cyber re-insurance for the highest layer of risk, whereas the state governments are responsible for private cyber-insurers to cover traditional cyber-risks. Thus, the primary force of such an act is that it will require a coordinated equilibrium in which all cyber-insurers must find it economically incentive compatible to offer cyber coverage upon a cyber-catastrophe; thereby, quite systematically (in accordance to results obtained by us in the article), government interventions to support catastrophe cyber-insurance markets can, in effect, boost the success of cyber re-insurance markets. However, in extreme circumstances (aggregate losses in the order of hundreds of billions or trillions of U.S. dollars), the market for cyber re-insurance can fail, and the governments need to take up the mantle of cyber-risk coverage. As a current related example in an analogous context of the recent 2020 COVID-19 outbreak, we observed that governments from most parts of the world, in order to boost the economy, provided monetary support (analogous to re-insurance) of amounts (in multi-billion U.S. dollars) far greater than any re-insurer is incentive compatible with or capable of. Our rationale behind promoting necessary government/regulatory intervention in the aftermath of catastrophic cyber-attacks stems from the fact that the COVID attack only causes disruptions of a certain section of the service sector,[12] leading to a negative commercial impact of multi-billion U.S. dollars, whereas a cyber-COVID-19 could cause a disruption in a major portion or the entire (IT/IoT-driven) service sector, consequently leading to an adverse commercial impact of trillions of U.S. dollars. Such amounts of risk can only be covered via government agencies, in addition to voluntary support by wealthy private agencies (e.g., Google, Facebook) playing integral roles in the IT/IoT services business.

9 ARTICLE SUMMARY

The IoT age will unlock over USD 10 trillion in new economic activity globally with a projected 50 billion IoT devices to be in operation by 2025. While many of these devices will work well and benefit society, many will be vulnerable even to naive cyber-attacks. To make things worse, due to common software platforms and applications running on many of these devices, the attack-impact

[12]People can still work online by staying at home, which might not be possible in case of a cyber-lockout.

effects will be correlated in time and space. Consequently, when these fail, it will not be unreasonable to expect significant aggregate cyber-loss impacts with statistically very heavy tails, simply because with increasing inter-connectivity, the cost of a failed connection goes higher. In this article, we investigated the robustness of cyber re-insurance mechanisms to cover aggregate cyber-losses incurred by organizations in a service-liability setting, under events of catastrophic cyber-attacks. Such attacks, often characterized by a specific form of heavy-tailed cyber-risk distributions, are increasingly (and also predictably) on the rise due to the emergence and pervasive prevalence of IoT technologies in myriad application spaces covering all walks of life. *Surprisingly, as a negative result for society in the event of such catastrophes, we proved via a game-theoretic analysis that it may not be economically incentive compatible (even under the ideal constraint of i.i.d. cyber-risk distributions) at a Nash equilibrium for limited liability-taking risk-averse cyber-insurance companies to offer cyber re-insurance solutions,* a sharp contrast to existing empirical results that show feasibility of cyber-risk spreading for (non-catastrophic) i.i.d. cyber-risks, this despite the existence of large enough market capacity to achieve full cyber-risk sharing. We validated this result via an experimental evaluation based on a real-world cyber-attack setting on a commercial cloud infrastructure. Our efforts serve as a mere conservative projection of larger-scale IoT-driven supply chain service networks. A potential failure to achieve a working cyber (re-)insurance market in such demanding situations strongly calls for centralized government regulatory action/intervention to promote cyber-risk sharing through re-insurance activities for the benefit of service-networked societies.

REFERENCES

Carlo Acerbi. 2002. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance* 26, 7 (2002), 1505–1518.

Franklin Allen and Douglas Gale. 2000. Financial contagion. *Journal of Political Economy* 108, 1 (2000), 1–33.

R. Anderson and T. Moore. 2009. Information security: Where computer science, economics and psychology meet. *Philosophical Transactions of the Royal Society* 367 (2009), 2717–2727.

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. 1999. Coherent measures of risk. *Mathematical Finance* 9, 3 (1999), 203–228.

Robert B. Ash, B. Robert, Catherine A. Doleans-Dade, and A. Catherine. 2000. *Probability and Measure Theory*. Academic Press.

Per Bak, Kan Chen, José Scheinkman, and Michael Woodford. 1993. Aggregate fluctuations from independent sectoral shocks: Self-organized criticality in a model of production and inventory dynamics. *Ricerche Economiche* 47, 1 (1993), 3–30.

Richard S. Betterley. 2017. Cyber/privacy insurance market survey. *The Betterley Report.*

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

Lawrence Blume, David Easley, Jon Kleinberg, Robert Kleinberg, and Éva Tardos. 2011. Which networks are least susceptible to cascading failures? In *IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS'11)*. IEEE, 393–402.

Rainer Böhme and Gaurav Kataria. 2006. Models and measures for correlation in cyber-insurance. In *WEIS*, Vol. 2. 3 pages.

Rainer Böhme, Stefan Laube, and Markus Riek. 2019. A fundamental approach to cyber risk analysis. *Variance* 12, 2 (2019), 161–185.

Rainer Böhme, Galina Schwartz, et al. 2010. Modeling cyber-insurance: Towards a unifying framework. In *WEIS*.

John Adrian Bondy, Uppaluri Siva Ramachandra Murty, et al. 1976. *Graph Theory with Applications*. Vol. 290. Macmillan London.

Jean-Philippe Bouchaud and Marc Potters. 2003. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*. Cambridge University Press.

Stephen Boyd and Lieven Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press.

Véronique Bruggeman, Michael G. Faure, and Karine Fiore. 2010. The government as reinsurer of catastrophe risks?*Geneva Papers on Risk and Insurance-Issues and Practice* 35, 3 (2010), 369–390.

Jacques Bughin, Michael Chui, and James Manyika. 2015. An executive's guide to the Internet of Things. *McKinsey Quarterly* 4 (2015), 92–101.

Nicholas G. Carr. 2003. IT doesn't matter. *Educause Review* 38 (2003), 24–38.
Robert L. Carter. 2013. *Reinsurance*. Springer Science & Business Media.
Krishna Chinthapalli. 2017. The hackers holding hospitals to ransom. *BMJ: British Medical Journal (Online)* 357 (2017).
Fan Chung and Linyuan Lu. 2002. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics* 6, 2 (2002), 125–145.
Andrew Coburn, Eireann Leverett, and Gordon Woo. 2018. *Solving Cyber Risk: Protecting Your Company and Society*. Wiley.
Wikipedia Contributors. 2018. 2007 cyberattacks on Estonia. *Wikipedia*.
John Cotter and Kevin Dowd. 2006. Extreme spectral risk measures: An application to futures clearinghouse margin requirements. *Journal of Banking & Finance* 30, 12 (2006), 3469–3485.
Savino Dambra, Leyla Bilge, and Davide Balzarotti. 2020. SoK: Cyber insurance–technical challenges and a system security roadmap. In *2020 IEEE Symposium on Security and Privacy (SP'20)*. 293–309.
Steven N. Durlauf. 1993. Nonergodic economic growth. *Review of Economic Studies* 60, 2 (1993), 349–366.
Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. 2016. Hype and heavy tails: A closer look at data breaches. *Journal of Cybersecurity* 2, 1 (2016), 3–14.
Martin Eling and Werner Schnell. 2020. *Extreme Cyber Risks and the Nondiversification Trap*. Technical Report. Working Paper University of St. Gallen. https://www.alexandria.unisg.
Martin Eling and Jan Wirfs. 2019. What are the actual costs of cyber risk events? *European Journal of Operational Research* 272, 3 (2019), 1109–1119.
Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. 2013. *Modelling Extremal Events: For Insurance and Finance*. Vol. 33. Springer Science & Business Media.
Paul Embrechts, Alexander McNeil, and Daniel Straumann. 2002. Correlation and dependence in risk management: Properties and pitfalls. *Risk Management: Value at Risk and Beyond* 1 (2002), 176–223.
Catherine Forbes, Merran Evans, Nicholas Hastings, and Brian Peacock. 2011. *Statistical Distributions*. John Wiley & Sons.
Xavier Freixas, Bruno M. Parigi, and Jean-Charles Rochet. 2000. Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of Money, Credit and Banking* (2000), 611–638.
Kenneth A. Froot, David S. Scharfstein, and Jeremy C. Stein. 1993. Risk management: Coordinating corporate investment and financing policies. *Journal of Finance* 48, 5 (1993), 1629–1658.
Prasanna Gai and Sujit Kapadia. 2010. Contagion in financial networks. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. The Royal Society, rspa20090410.
Ayalvadi Ganesh, Laurent Massoulié, and Don Towsley. 2005. The effect of network topology on the spread of epidemics. In *Proceedings of the IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'05)*., Vol. 2. IEEE, 1455–1466.
Mark Granovetter. 1978. Threshold models of collective behavior. *American Journal of Sociology* 83, 6 (1978), 1420–1443.
Andy Greenberg. 2017. How an entire nation became Russia's test lab for cyberwar.
Andy Greenberg. 2018. The Untold Story of NotPetya, the Most Devastating Cyberattack in History. Vol. 22.
Steve Grobman. 2018. When nation-states hack the private sector for intellectual property. *The Hill*.
Kat Hall. 2017. UK hospital meltdown after ransomware worm uses NSA vuln to raid IT.
Hemantha Herath and Tejaswini Herath. 2011. Copula-based actuarial model for pricing cyber-insurance policies. *Insurance Markets and Companies: Analyses and Actuarial Computations* 2, 1 (2011), 7–20.
Annette Hoffman. 2007. Internalizing externalities of loss prevention through insurance monopoly. *Geneva Risk and Insurance Review* 32, 1 (2007), 91–111.
Glyn A. Holton. 2003. *Value-at-Risk*. Academic Press.
Rustam Ibragimov. 2009. Portfolio diversification and value at risk under thick-tailedness. *Quantitative Finance* 9, 5 (2009), 565–580.
Rustam Ibragimov, Dwight Jaffee, and Johan Walden. 2009. Nondiversification traps in catastrophe insurance markets. *Review of Financial Studies* 22, 3 (2009), 959–993.
Risk Management Solutions Inc. 2016. Managing Cyber Insurance Accumulation Risk.
B. Johnson, R. Bohme, and J. Grossklags. 2011. Security games with market insurance. In *GameSec*.
Denis Kessler. 2014. Why (re) insurance is not systemic. *Journal of Risk and Insurance* 81, 3 (2014), 477–488.
Mohammad Mahdi Khalili, Parinaz Naghizadeh, and Mingyan Liu. 2018. Designing cyber insurance policies: The role of pre-screening and security interdependence. *IEEE Transactions on Information Forensics and Security* 13, 9 (2018), 2226–2239.
Jocelyn Krystlik. 2017. With GDPR, preparation is everything. *Computer Fraud & Security* 2017, 6 (2017), 5–8.
Richard J. La. 2016. Interdependent security with strategic agents and cascades of infection. *IEEE/ACM Transactions on Networking* 24, 3 (2016), 1378–1391.
Richard La. 2018a. Influence of clustering on cascading failures in interdependent systems. *IEEE Transactions on Network Science and Engineering* 6, 3 (2018), 351–363.

Richard J. La. 2018b. Cascading failures in interdependent systems: Impact of degree variability and dependence. *IEEE Transactions on Network Science and Engineering* 5, 2 (2018), 127–140.

M. Lelarge and J. Bolot. 2009a. Economic incentives to increase security in the Internet: The case for insurance. In *IEEE INFOCOM*.

Marc Lelarge and Jean Bolot. 2009b. Economic incentives to increase security in the internet: The case for insurance. In *IEEE INFOCOM 2009*. IEEE, 1494–1502.

Mark Li. 2018. SCOR paper. (2018).

Jan Lorenz, Stefano Battiston, and Frank Schweitzer. 2009. Systemic risk in a unifying framework for cascading processes on networks. *European Physical Journal B* 71, 4 (2009), 441.

Thomas Maillart and Didier Sornette. 2010. Heavy-tailed distribution of cyber-risks. *European Physical Journal B* 75, 3 (2010), 357–364.

Albert W. Marshall, Ingram Olkin, et al. 1974. Majorization in multivariate distributions. *Annals of Statistics* 2, 6 (1974), 1189–1200.

Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. 1979. *Inequalities: Theory of Majorization and Its Applications*. Vol. 143. Springer.

Solan Eilan Maschler Michael, and Zamir Shmuel. 2013. *Game Theory*, Vol. 979. pp. xxvi. Cambridge University Press.

Alexander J. McNeil, Rüdiger Frey, Paul Embrechts, et al. 2015. Quantitative risk management: Concepts. *Economics Books* (2015).

Pascal Millaire. 2016. 3 reasons why the insurance industry will never be the same after the Mirai DDoS attack.

Michael Molloy and Bruce Reed. 1995. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms* 6, 2–3 (1995), 161–180.

Michael Molloy and Bruce Reed. 1998. The size of the giant component of a random graph with a given degree sequence. *Combinatorics, Probability and Computing* 7, 3 (1998), 295–305.

Vincenzo Morabito. 2017. The security of blockchain systems. In *Business Innovation through Blockchain*. Springer, 61–78.

Stephen Morris. 2000. Contagion. *Review of Economic Studies* 67, 1 (2000), 57–78.

Arunabha Mukhopadhyay, Samir Chatterjee, Debashis Saha, Ambuj Mahanti, and Samir K. Sadhukhan. 2013. Cyber-risk decision models: To insure IT or not? *Decision Support Systems* 56 (2013), 11–26.

Parinaz Naghizadeh and Mingyan Liu. 2014. Voluntary participation in cyber-insurance markets. In *Workshop on the Economics of Information Security (WEIS'14)*.

Parinaz Naghizadeh and Mingyan Liu. 2016. Exit equilibrium: Towards understanding voluntary participation in security games. In *The 35th Annual IEEE International Conference on Computer Communications (IEEE INFOCOM'16)*. IEEE, 1–9.

Mark E. J. Newman, Steven H. Strogatz, and Duncan J. Watts. 2001. Random graphs with arbitrary degree distributions and their applications. *Physical Review E* 64, 2 (2001), 026118.

Oracle. 2019. The impact of emerging technology on CX excellence. *Oracle Report*.

Ranjan Pal and Leana Golubchik. 2010. Analyzing self-defense investments in internet security under cyber-insurance coverage. In *2010 IEEE 30th International Conference on Distributed Computing Systems*. IEEE, 339–347.

Ranjan Pal, Leana Golubchik, and Konstantinos Psounis. 2011. Aegis a novel cyber-insurance model. In *International Conference on Decision and Game Theory for Security*. Springer, 131–150.

Ranjan Pal, Leana Golubchik, Konstantions Psounis, and Tathagata Bandyopadhyay. 2019. On robust estimates of correlated risk in cyber-insured IT firms: A first look at optimal AI-based estimates under "Small" data. *ACM Transactions on Management Information Systems (TMIS)* 10, 3 (2019), 1–18.

Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. 2014. Will cyber-insurance improve network security? A market analysis. In *2014 Proceedings IEEE (INFOCOM'14)*. IEEE, 235–243.

Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. 2017. Security pricing as enabler of cyber-insurance a first look at differentiated pricing markets. *IEEE Transactions on Dependable and Secure Computing* 16, 2 (2017), 358–372.

Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Pan Hui. 2018. Improving cyber-security via profitable insurance markets. *ACM SIGMETRICS Performance Evaluation Review* 45, 4 (2018), 7–15.

Ranjan Pal, Ziyuan Huang, Xinlong Yin, Mingyan Liu, Sergey Lototsky, and Jon Crowcroft. 2020a. Sustainable catastrophic cyber-risk management in IoT societies. In *Appeared in 2010 Winter Simulation Conference*. IEEE/INFORMS.

Ranjan Pal, Ziyuan Huang, Xinlong Yin, Mingyan Liu, Sergey Lototsky, and Jon Crowcroft. 2020b. Sustainable catastrophic cyber-risk management in IoT societies. In *IEEE/ACM/INFORMS Winter Simulation Conference*.

Ranjan Pal, Ziyuan Huang, Xinlong Yin, Sergey Lototsky, Swades De, Sasu Tarkoma, Mingyan Liu, Jon Crowcroft, and Nishanth Sastry. 2020c. Aggregate cyber-risk management in the IoT age: Cautionary statistics for (re) insurers and likes. *IEEE Internet of Things Journal* (2020).

Ranjan Pal, Konstantinos Psounis, Jon Crowcroft, Frank Kelly, Pan Hui, Sasu Tarkoma, Abhishek Kumar, John Kelly, Aritra Chatterjee, Leana Golubchik, et al. 2020d. When are cyber blackouts in modern service networks likely? A network oblivious theory on cyber (re) insurance feasibility. *ACM Transactions on Management Information Systems (TMIS)* 11, 2 (2020), 1–38.

Chen Peng, Maochao Xu, Shouhuai Xu, and Taizhong Hu. 2018. Modeling multivariate cybersecurity risks. *Journal of Applied Statistics* 45, 15 (2018), 2718–2740.

Frank Proschan. 1965. Peakedness of distributions of convex combinations. *Annals of Mathematical Statistics* 36, 6 (1965), 1703–1706.

Alison DeNisco Rayome. 2017. 33% of businesses hit by DDoS attack in 2017, double that of 2016. *TechRepublic* (Nov. 2017). Retrieved from https://www.techrepublic.com/article/33-of-businesses-hit-by-ddos-attack-in-2017-double-that-of-2016/.

Markus Riek and Rainer Böhme. 2018. The costs of consumer-facing cybercrime: An empirical exploration of measurement issues and estimates. *Journal of Cybersecurity* 4, 1 (2018), tyy004.

Jordan Robertson and Michael Riley. 2018. The big hack: How china used a tiny chip to infiltrate us companies. *Bloomberg Businessweek*, vol. 4.

S. Romanovsky. 2013. *Comments to the Department of Commerce on Incentives to Adopt Improved Cyber-Security Practices*. April 2013.

Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O'Flynn. 2017. IoT goes nuclear: Creating a ZigBee chain reaction. In *IEEE Symposium on Security and Privacy (SP'17)*. IEEE, 195–212.

Haskell P. Rosenthal. 1973. On subspaces of Lp. *Annals of Mathematics* (1973), 344–373.

Sheldon M. Ross. 2014. *Introduction to Probability Models*. Academic Press.

Michael Rothschild and Joseph E. Stiglitz. 1970. Increasing risk: I. A definition. *Journal of Economic Theory* 2, 3 (1970), 225–243.

Xinyi Wu and Gary Gereffi. 2018. Amazon and Alibaba: Internet governance, business models, and internationalization strategies. In *International Business in the Information and Digital Age*. Emerald Publishing Limited.

Paul A. Samuelson. 1967. General proof that diversification pays. *Journal of Financial and Quantitative Analysis* 2, 1 (1967), 1–13.

Nikhil Shetty, Galina Schwartz, Mark Felegyhazi, and Jean Walrand. 2010. Competitive cyber-insurance and internet security. In *Economics of Information Security and Privacy*. Springer, 229–247.

Thom Tracy. 2016. Apple Stock: Analyzing 5 Key Customers (AAPL). Investopedia.

Vladimir V. Uchaikin and Vladimir M. Zolotarev. 2011. *Chance and Stability: Stable Distributions and Their Applications*. Walter de Gruyter.

Duncan J. Watts. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences* 99, 9 (2002), 5766–5771.

Jonathan William Welburn and Aaron Strong. 2019. *Systemic Cyber Risk and Aggregate Impacts*. RAND.

Spencer Wheatley, Thomas Maillart, and Didier Sornette. 2016. The extreme risk of personal data breaches and the erosion of privacy. *European Physical Journal B* 89, 1 (2016), 1–12.

Zack Whittaker. 2016. Mirai botnet attack hits thousands of home routers, throwing users offline. *ZDNet 29*.

Benjamin Wootton. 2017. Who's using Amazon web services?*Contingo*.

Maochao Xu and Lei Hua. 2019. Cybersecurity insurance: Modeling and pricing. *North American Actuarial Journal* 23, 2 (2019), 220–249.

Maochao Xu, Lei Hua, and Shouhuai Xu. 2017. A vine copula model for predicting the effectiveness of cyber defense early-warning. *Technometrics* 59, 4 (2017), 508–520.

Maochao Xu, Kristin M. Schweitzer, Raymond M. Bateman, and Shouhuai Xu. 2018. Modeling and predicting cyber hacking breaches. *IEEE Transactions on Information Forensics and Security* 13, 11 (2018), 2856–2871.

Vladimir M. Zolotarev. 1986. *One-Dimensional Stable Distributions*. Vol. 65. American Mathematical Society.

Received March 2020; revised August 2020; accepted December 2020