When Are Cyber Blackouts in Modern Service Networks Likely?: A Network Oblivious Theory on Cyber (Re)Insurance Feasibility

Pal, Ranjan

2020-07

Pal , R , Psounis , K , Crowcroft , J , Kelly , F , Hui , P , Tarkoma , S , Kumar , A , Kelly , J , Chatterjee , A , Golubchik , L , Sastry , N & Nag , B 2020 , ' When Are Cyber Blackouts in Modern Service Networks Likely?: A Network Oblivious Theory on Cyber (Re)Insurance Feasibility ' , ACM transactions on management information systems , vol. 11 , no. 2 , 5 . https://doi.org/10.1145/338

http://hdl.handle.net/10138/318422
https://doi.org/10.1145/3386159

acceptedVersion

When Are Cyber Blackouts in Modern Service Networks Likely?: A Network Oblivious Theory on Cyber (Re)Insurance Feasibility

RANJAN PAL, University of Michigan Ann Arbor, USA, University of Cambridge, UK
KONSTANTINOS PSOUNIS, University of Southern California, USA
JON CROWCROFT and FRANK KELLY, University of Cambridge, UK
PAN HUI, SASU TARKOMA, and ABHISHEK KUMAR, University of Helsinki, Finland
JOHN KELLY and ARITRA CHATTERJEE, Envelop Risk
LEANA GOLUBCHIK, University of Southern California, USA
NISHANTH SASTRY, King's College London, UK
BODHIBRATA NAG, Indian Institute of Management Calcutta, India

Service liability interconnections among globally networked IT- and IoT-driven service organizations create potential channels for cascading service disruptions worth billions of dollars, due to modern cyber-crimes such as DDoS, APT, and ransomware attacks. A natural question that arises in this context is: *What is the likelihood of a cyber-blackout?*, where the latter term is defined as the probability that all (or a major subset of) organizations in a service chain become dysfunctional in a certain manner due to a cyber-attack at some or all points in the chain. The answer to this question has major implications to risk management businesses such as cyber-insurance when it comes to designing policies by risk-averse insurers for providing coverage to clients in the aftermath of such catastrophic network events. In this article, we investigate this question in general as a function of service chain networks and different cyber-loss distribution types. We show somewhat surprisingly (and discuss the potential practical implications) that, following a cyber-attack, the effect of (a) a network interconnection topology and (b) a wide range of loss distributions on the probability of a cyber-blackout and the increase in total service-related monetary losses across all organizations are *mostly* very small. The primary rationale behind these results are attributed to degrees of heterogeneity in the revenue base among organizations and the *Increasing Failure Rate* property of popular (i.i.d/non-i.i.d) loss distributions, i.e., log-concave cyber-loss distributions. The result will enable risk-averse cyber-risk managers to safely infer the impact of cyber-attacks in a worst-case network and distribution oblivious setting.

Authors' addresses: R. Pal and J. Crowcroft, Computer Laboratory, University of Cambridge, Cambridge, CB3 OFD, UK; emails: {rp631, jac22}@cam.ac.uk; K. Psounis and L. Golubchik, Department of Computer Science and Electrical Engineering, University of Southern California, Los Angeles, California, 90089, USA; emails: {kpsounis, leana}@usc.edu; A. Kumar and S. Tarkoma, Department of Computer Science, University of Helsinki, Yliopistonkatu 4, 00100 Helsinki, Finland; emails: {abhishek.kumar, sasu.tarkoma}@helsinki.fi; P. Hui, Department of Computer Science, University of Helsinki, Yliopistonkatu 4, 00100 Helsinki, Finland and Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong; email: pan.hui@cs.helsinki.fi; F. Kelly, Statistical Laboratory, University of Cambridge, Cambridge, CB3 OFD, UK; email: F.P. Kelly@statslab.cam.ac.uk; J. Kelly, Envelop Risk, Washington D.C.; email: john.kelly@enveloprisk.com; A. Chatterjee, Envelop Risk, Bermuda; email: aritra81@gmail.com; N. Sastry, King's College London, London WC2R 2LS, UK; email: nishanth.sastry@kcl.ac.uk; B. Nag, Indian Institute of Management Calcutta, Calcutta 700104, India; email: bnag@iimcal.ac.in.

2158-656X/2020/05-ART5 $15.00
https://doi.org/10.1145/3386159

CCS Concepts: • **Mathematics of computing → Probability and statistics**; **Multivariate statistics**; • **Security and privacy → Human and societal aspects of security and privacy**;

Additional Key Words and Phrases: Service network, cyber-blackout, systemic risk

ACM Reference format:
Ranjan Pal, Konstantinos Psounis, Jon Crowcroft, Frank Kelly, Pan Hui, Sasu Tarkoma, Abhishek Kumar, John Kelly, Aritra Chatterjee, Leana Golubchik, Nishanth Sastry, and Bodhibrata Nag. 2020. *When Are Cyber Blackouts in Modern Service Networks Likely?*: A Network Oblivious Theory on Cyber (Re)Insurance Feasibility. *ACM Trans. Manage. Inf. Syst.* 11, 2, Article 5 (May 2020), 38 pages.
https://doi.org/10.1145/3386159

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, and the like, 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., a 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. As well-documented commercial cyber-attack examples in reality, the very recent *Mirai* DDoS (2016), *NotPetya* ransomware (2017), and *WannaCry* ransomware (2017) attacks caused havoc among firms in various industries across the globe, resulting in huge financial losses for the firms due to their being deemed dysfunctional in providing service to customers.

1.1 Research Motivation

In the wake of major targeted corporate cyber-attacks (e.g., attacks on Sony, Target) 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 (*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. However, market surveys suggest that demand for cyber-insurance significantly exceeds the capacity currently provided by the insurance industry. The primary reason that most insurers give for being cautious about expanding capacity is the accumulation risk posed by cyber-threats. The main fear among insurers here is that cyber-threats are inherently scalable and systemic through their spread via network interconnectivity—a single malicious email generated by a botnet activity as part of a social engineering attack can result in an entire organization becoming dysfunctional with respect to the service it provides, and in turn potentially affecting 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 on for providing services, it is quite likely that the insurance company of a certain root organization would need to bear the responsibility of covering a huge aggregate/accumulated risk of all or multiple organizations in the service chain [Millaire 2016]. Shouldering this responsibility clearly may not be aligned with satisfying the budget constraints and profit requirements of most commercial risk-averse cyber-insurers, leave alone risk-tracking

and risk-data availability challenges they might need to overcome to implement accumulative coverage policies [Millaire 2016]. The above argument is still an unproven skepticism and the current state of mind of cyber-insurance suppliers who are extremely risk-averse (primarily by perception) and do not open up their complete coverage capacity to the client side, even if they have spare coverage capacity. Our initial feel/conjecture is that such skepticism is unwanted for in the timely necessary interests of the opening up of cyber-insurance markets, specially in the wake of increasingly complex cyber-attacks, and is creating a big bottleneck in leveraging the benefits of complete coverage capacity of insurers to their clients. In this paper we wish to conduct a thorough mathematical investigation to judge the validity of such a skepticism.

Our Focus - Given a service chain network, our focus in this paper is to estimate in a manner oblivious to a network and cyber-loss distribution in the worst case, the probability that *all* or a *major subset* of organizations (network nodes) in the network become dysfunctional in a certain manner (e.g., unable to provide cloud connectivity, inability to protect customer privacy, disruption of energy services) to provide service in the event of a cyber-attack, a situation which we define as a *cyber-blackout* (See Section 4.1 for details]. The requirement of a robust estimation environment that is oblivious to network and cyber-loss distributions can often arise for a large service liability network setting, where cyber-risk managers may not have complete knowledge about the topology and statistical cyber-attack distributions. A robust estimate of the probability of a cyber-blackout is a necessary pre-requisite for considering the expansion of the service capacity of risk management products such as cyber-insurance. In scenarios of cascading cyber-risks, the probability bounds will act as a valuable input to cyber-insurance firms to allocate optimal portfolios among insurance and re-insurance investments. In addition to the above, we will also investigate the practical implications of the likelihood and scale of cyber-blackouts on cyber-insurance ecosystems of today and the near future.

1.2 Research Contributions

We make the following research contributions in this article.

- —We design a graph-based model of service obligations, *GSOM*, between organizations in a service chain network. Our model specifies a set of nodes that represent service organizations together with the edges that represent service liability relationships between them. In the event of a cyber-attack, given the values of losses (either deterministic or stochastic) at the nodes in the network, GSOM computes via solving a fixed-point problem, the vector of service valuations that clears the network, and identifies the nodes in the chain that are dysfunctional to provide service. GSOM is very useful for analyzing how service-related losses propagate through an organizational service chain (see Section 3).
- —Using GSOM, given the joint distribution of service-related losses across the network nodes (organizations) in the event of a cyber-attack, we analyze the probability of contagion that target organizations become dysfunctional due to a given organization somewhere in the network becoming dysfunctional. In this regard, we answer two important questions: (i) how likely it is that a given set of target organizations will become dysfunctional due to contagion from a *single source* organization, as compared to the likelihood that they become dysfunctional from direct losses to their own service-related assets that does not require dependency on other nodes?, and (ii) how much does the underlying network of service dependencies contribute to the *increase* in (a) the probability of dysfunction of target nodes and (b) the corresponding expected value of losses, compared to a situation when there are no network links, i.e., each organization completely relies on its own resources to provide customer service? Our analysis is very useful for analyzing the chance of a cyber-blackout event, and the effect of the underlying liability network topology and cyber-risk distribution

on cascading node dysfunction parameters. As part of our results, we derive a general formula that surprisingly shows that the probability of a cyber-blackout is larger *mostly* in the absence of network connectivity than that in the presence of network connectivity, implying that simple network spillover effects have a limited impact with respect to service obligations between heterogeneous (in terms of monetary assets) organizations. We also show that network spillover effects are surprisingly small mostly under a wide range of joint distributions for plausible values of model parameters, regardless of the service dependency network topology. (see Section 4 for details) - the rationale behind these results being attributed to degrees of heterogeneity in organizational revenue bases, and the Increasing Failure Rate (IFR) property of popular (iid/non-iid) cyber-loss distributions (e.g., log-concave distributions).

—We expand the set of cyber-attack sources from a single node to multiple nodes, and study the negative impact of simultaneous attacks on the entire network of organizations. Under a wide range of general (iid/non-iid) loss distributions, we again surprisingly show that the increase in total (summed over all nodes in the organizational network) value of service-related losses due to network interconnections are mostly small, regardless of the network structure (see Section 5 for details).

Our results will enable risk-averse cyber-risk managers to safely infer the impact of cyber-attacks in a worst-case network and distribution oblivious setting.

2 MOTIVATION EXAMPLES FOR CYBER-BLACKOUT STUDY

In this section, we provide motivational background for our article by describing various well-known attack scenarios that are capable of launching a cyber-blackout, thereby potentially presenting an accumulative coverage setting for a cyber-insurance provider. We describe coverage accumulation scenarios for *seven* key example processes [Inc. 2016] of cyber-loss in today's digital age. The examples highlight how correlated cyber-losses could impact a portfolio of cyber-insurance policies, and peep into the rationale of how a large number of accounts/organizations might suffer systemic losses from a single underlying cause.

Cyber-Data Exfiltration. This process relates to the systemic release of confidential customer records from many corporate enterprises (organizations). Some of the highest profile cyber-incidents (e.g., the *Sony*, *Target*, and *Equifax* cyber-incidents [Risk Management Solutions, Inc. 2016]) have been data breaches[1]: the loss of confidential data from organizations that breach the privacy of their customers, employees, clients, or counterparts. This has proved costly to the enterprise, resulting in notification costs, credit monitoring services, and compensation pay-outs to all the individuals/organizations whose data was compromised, together with regulatory fines, response and forensic costs, and sometimes substantial litigation costs. The total accumulative losses to data breaches (both, first-party, and third-party losses faced from organizations in the service chain complaining of privacy breach of their data) faced by individual organizations have been instrumental in driving the expansion of the cyber-insurance market, as companies seek protection and risk partners in helping with response services.

Another burning example of cyber-data exfiltration might arise from the recently operative General Data Protection Regulation (GDPR) [Krystlik 2017]. The key theme of GDPR, operative in the EU from May 25th, is that each of us owns our own data. Any company (EU local/EU multinational/companies worldwide operating with data of EU subjects including residents, citizens, and

[1]Types of data include personal identity information (PII), payment and credit card information (PCI), protected health information (PHI), commercial confidential information (CCI), and intellectual property (IP).

tourists) must therefore explicitly request permission to use any of our personal data, explaining why it would like to do so, and for how long. If we so agree, we can later withdraw our permission at any time. All of these rights must be provided to us by each company free of charge. One consequence is that each company must know, and (dynamically) document, what information (if any) they have about each individual. This may be a particular challenge for large, established corporations, since data about individuals may be spread across different business units and multiple databases, spreadsheets, off-site backup copies, or even paper archives, that would make synchronous dynamic updating of data difficult. Thus, will open up avenues for cyber-data exfiltration thereby leading to the aforementioned situation of accumulative losses due to a data breach.

Denial-of-Service Attack. This process relates to attacks that disable websites and disrupt online business activity across multiple organizations. Denial-of-Service (DoS) and Distributed DoS attacks are common methods of disrupting website business activities by bombarding them with traffic (e.g., the *Mirai* botnet-induced DDoS attack [Whittaker 2016]). According to a Kaspersky Lab report, 33% of organizations experienced a DoS attack on their websites in the past few years, and one in eight of those attacks overwhelmed website resilience and rendered Internet services unavailable [DeNisco Rayome 2017]. In April 2007, following a diplomatic row with Russia over a Soviet war memorial, Estonia was subject to DDos attacks which caused temporary shut down of infrastructure including everything from online banking and mobile phone networks to government services and access to health care information [Wikipedia Contributors 2018]. For a given organization, the cost of the business interruption caused by a DDoS attack of any particular duration is determined by the Internet dependency of the insured company, i.e., the amount of revenue that would be lost per hour of Internet failure or connectivity loss. The capture of this information makes it possible to assess the accumulative loss of revenue that a given insured organization may be liable for (due to organizational dependencies in a service chain), from the potential for Internet outage in general. As another line of recent target applications for DDoS attacks is the catastrophic disruption of critical infrastructure services in the electricity, manufacturing, and transportation sectors by APT-driven DDoS attack vectors. A well-known instance of such an attack type led to a series of power blackouts in Ukraine [Greenberg 2017] in the last few years that caused significant damage to people's lives and business activities. Such scenarios also lead to accumulative coverage for cyber-insurers.

Cloud Service Provider Failure. This process relates to the scenario when large number of organizations have business operations disrupted by losing cloud-based functionality in the event of a major cloud service provider (CSP) suffering a service disruption due to a cyber-attack. The digital economy is increasingly dependent on cloud services and a rapidly growing number of companies make use of a CSP by outsourcing elements of their data storage, analytics, and information technology functions. If a CSP were to fail (e.g., *AWS* outage (2011), *Gmail* outage (2010), *Microsoft Sidekick* outage (2009) [Risk Management Solutions 2016]), then their customers would suffer business losses and hold the CSP liable for the loss. A CSP failure could also be the source of the exfiltration of confidential data records or claims for data and software loss if data files were irrevocably deleted. This provides an accumulation issue for cyber-insurance where there is potential for a large number of organizations (and their subsequent business clients in a service chain) to make a claim for business interruption if a major provider of cloud services were to have a lengthy outage or failure, from any cause. The systemic dimension of cyber-risk concerns the triggering of large numbers of claims from companies that are CSP customers. The customers and their insurers may attempt to recover their loss payouts from the CSP (and the insurer of the CSP).

Compromise of Financial Transactions. This process relates to theft of large sums in cyber-attacks on multiple enterprises (organizations) that carry out financial transactions. Insurers

offer coverage to the financial services sector to cover losses that they might suffer from cyber-attacks, or computer based fraud, theft, or disruption occurring from compromising payment systems or technologies for managing financial transactions. Criminals have always targeted the money held in financial institutions, and physical bank robbery has given way to cyber-crime as the preferred technique (e.g., *Carbanak APT Attack* (2013–2015), *Drinkman and Kalinin Attack* (2013) [Risk Management Solutions, Inc. 2016]). Although a very large number of companies of all different types carry out financial transactions, ranging from retail to e-commerce, the transaction systems that carry the financial flows are the specific liabilities of financial transaction companies. The potential for widespread and systemic claims across all the different sectors of the economy from subverting payments after the point of sale are constrained by the legal liabilities being confined to the financial services companies operating the payment transfers. Thus, transaction risk is mostly aggregated in banking and payment management companies and investment management systems.

Cyber-Extortion through Ransomware. This process relates to the event when many companies are held to ransom by payoff seeking hackers disabling IT functionality. Cyber-extortion is a rapidly growing area of organized cyber-crime using ransomware—malicious software that locks up data or disrupts business until companies make a payoff. This has been a common method of extorting individuals and small businesses for some years (e.g., *LA Children's Presbyterian Hospital Attack* (2016) [Chinthapalli 2017], *Bitfinex Attack* (2017) [Morabito 2017]). Cyber-criminals are increasingly scaling up their operations and using extortion more commonly against larger companies as they gain confidence and technical expertise. In 2017, UK hospitals effectively shut down and had to turn away non-emergency patients after *WannaCrypt* ransomware ransacked its networks [Hall 2017]. In the same year, Maersk, the world's largest container shipping company, was hit by *NotPetya* ransomware attack [Greenberg 2018]. Although ransomware that encrypts data and locks computers is the most common type of extortion, companies may also be asked to make payoffs to avert the threat of other cyber-attack types including denial-of-service attacks, data exfiltration breach, and sabotage to deny a company Internet or cloud services. Insurance repayment for extortion is a common coverage in many standalone affirmative cyber-liability products in the market, and around three-quarters of products offer this. Following from the previously mentioned process examples, accumulative risk is something a cyber-insurance company needs to deal with in the case of extortions.

Aggregate Losses due to Cyber-Warfare. Highly untraceable acts of modern cyber-warfare or cyber-terrorism by nation states aimed to achieve political and corporate gains can lead to aggregate losses incurred by organizations. On this note, as per a recent report by *Bloomberg Businessweek* [Robertson and Riley 2018], data center equipment run by Amazon Web Services and Apple may have been subject to surveillance from the Chinese government via a tiny and virtually undetectable microchip inserted during the equipment manufacturing process. These illicit microchips were capable of instructing the device in which they were embedded to communicate with unauthorized computers located elsewhere on the Internet and preparing the device's operating system to accept new code, and hence, enabling attackers to alter how the device functioned, however they wanted. As an example, attackers could use this to steal intellectual property (IP) of organizations and their service-providing clients, resulting in a situation of aggregate information leak targeted at the host [Grobman 2018]. For the microchip case, Netflix (Entertainment), BBC (News Broadcasting), Capital One Financial Corporation (Finance), Twitter (Social Media), and various departments of the US government were clients of Amazon Web Services [Wootton 2017; AWS Sales 2018]. Similarly, Best Buy (Consumer Electronics), Verizon Communications (Telecommunications), AT&T (Telecommunications), Sprint (Telecommunications), T-Mobile U.S.

(Telecommunications), were clients of Apple [Tracy 2016]. Here, Amazon and Apple may be held liable by their clients for IP loss damages inflicted during such attacks, thereby contributing to accumulative risk for a cyber-insurance company.

Aggregated Risk in IoT-Driven Smart Cities. In the near future, people are likely to populate their homes, offices, and neighborhood with a dense network of potentially billions of tiny transmitters and receivers which have ad-hoc networking abilities. These IoT devices can directly communicate amongst themselves, creating a new unintended communication medium that completely bypasses the traditional norms of communications such as telephony and the Internet. In a recent work, Ronen et al. [2017] has successfully demonstrated that even though IoT devices might be manufactured by popular and reputed firms deploying industry-standard cryptographic techniques, they can be still misused by hackers to spread infectious malware from one IoT device to all physically adjacent neighbors, causing city-wide disruptions which are very difficult to stop and investigate Ronen et al. [2017]. In the case of "city" insurance agencies insuring their clients in the future, aggregate cascading risks due to unavailability of service is something they might have to deal with.

3 SYSTEM MODEL

In this section, we propose our graph-based model of service obligations, GSOM, between organizations in a service chain network that will be used in this article to investigate and analyze cyber-blackout probabilities.

3.1 Basic Ingredients of GSOM

GSOM has four basic ingredients: (i) a set of n nodes $N = \{1, 2, \ldots, n\}$ characterizing organizations, (ii) an $n \times n$ liability matrix $\bar{P} = (\bar{p}_{ij})$ where $\bar{p}_{ij} \geq 0$ is the payment *due* from node i to node j in the event of a claim made by j on i in the aftermath of a cyber-attack (e.g., an organization claiming that due to CSP failure, it incurred a business loss worth a certain monetary amount) with $\bar{p}_{ii} = 0$, (iii) $\vec{c} = (c_1, c_2, \ldots, c_n) \in \mathrm{R}_+^n$, representing the vector of wealth/resource amount held by each node $i \in N$, that is not yet subject to a cyber-attack, and (iv) $\vec{b} = (b_1, b_2, \ldots, b_n) \in \mathrm{R}_+^n$, representing the vector of liability-free losses accrued by each node $i \in N$, in the aftermath of a cyber-attack. We make the general assumption in this article that organizational claims, wealth, and losses can be expressed monetarily in the event of a cyber-attack. *Also note that the liability matrix embeds the service chain network.* For each node $i \in N$, the following relationship holds:

$$w_i = c_i + \sum_{j \neq i} \bar{p_{ji}} - \bar{p}_i, \tag{1}$$

where w_i is the net wealth of node i in the aftermath of a cyber-incident (given that the claim payouts are appropriately meted out), and is unrestricted in sign, and $\bar{p}_i$ is the net liability of i. A negative value of w_i denotes the inability of organization i to pay out claims made by organizations liable on i. Observe that the net liability of i is expressed as

$$\bar{p}_i = b_i + \sum_{j \neq i} \bar{p}_{ij}.$$

Similarly, the net non-liability (assets) of organization i in the aftermath of a cyber-attack is given by $c_i + \sum_{j \neq i} \bar{p}_{ji}$. We illustrate the basic ingredients of the GSOM model via an example in the Appendix.

3.2 GSOM for Post-Attack Scene

Having discussed the basic elements of GSOM, our primary goal here is to build GSOM to handle the case when resources that have not yet been hit by a cyber-attack are suddenly subject to a loss that might trigger service disruption in a service chain network.

Let the amount c_i for each node i be subject to a random shock or loss of value X_i in the event of a cyber-attack, where X_i is a random variable taking values in the interval $[0, c_i]$. Thus, in the aftermath of the attack, resource amount c_i for node i is reduced to $c_i - x_i$, where x_i is the instance of X_i. Let $F(x_1, x_2, \ldots, x_n)$ be the joint cumulative function of these losses, that is central to analyzing the process of the spread of "organizational dysfunctionality" due to cyber-attacks. We note that it is important that a *necessary* (but not sufficient and complete) component to estimating or approximating F is the use of techniques like Monte Carlo simulation, percolation theory, or statistical mean field models, that popularly capture the spread of the infection (attack) vector (e.g., a virus, worm, bot) across a network, and is not the focus of our article. The interested reader is referred to Lelarge and Bolot [2009], Lorenz et al. [2009], Ganesh et al. [2005], and Gao et al. [2012] to get insights about some ways to mathematically evaluate this necessary component contributing to the value of F. *In our work, we adopt a conservative (and hence more challenging) approach of assuming general continuous forms of F for the purpose of analysis*, without focusing our efforts (via the use of the aforementioned necessary component) on finding/assuming *specific* continuous forms of F that might be setting-dependent.

Define the relative liabilities matrix $A = (a_{ij})$ to be the $n \times n$ matrix with the entries:

$$a_{ij} = \begin{cases} \frac{\bar{p}_{ij}}{\bar{p}_i}, \text{ if } \bar{p}_i > 0 \\ \\ 0, \text{if } \bar{p}_i = 0. \end{cases}$$

Thus, a_{ij} is the proportion of organization i's monetary obligations owed to organizations j in the aftermath of cyber-attack. Here, $a_{ij} \leq 1$ for each i and subsequently matrix A is substochastic.

Given a loss realization vector $\vec{x} = (x_1, \ldots, x_n) \geq 0$, our aim is to evaluate a vector that corresponds to the payments that balance monetary assets and liabilities at each node (organization). Based on the values in this vector, we will know whether an organization does have the ability to handle liabilities from other organizations in the service chain, in the event the latter hold the former liable for service disruptions due to a cyber-attack. In this article, we term the vector as a *clearing vector* due to its relevance in balancing assets and liabilities. Mathematically, we represent the clearing vector as $\vec{p}(\vec{x}) = \{p_i(\vec{x})\}; \vec{p}(\vec{x}) \in \mathrm{R}_+^n$, and it is evaluated as the solution to the following fixed-point equation:

$$p_i(\vec{x}) = \bar{p}_i \wedge \left(\sum_j p_j(\vec{x}) a_{ji} + c_i - x_i \right)_+, \tag{2}$$

where the structure $(\cdot)_+$ above indicates that if the value inside the parenthesis is less than zero, the value is zero. $\wedge$ denotes the min operator, where $\vec{x} \wedge \vec{y} = (\min[x_i, y_i], \ldots, \min[x_n, y_n])$ The solution to this equation, for each node i, is evaluated under a *pro-rata* allocation mechanism, i.e., the amount of unresolved liabilities at node i (when its net assets are less than its net liability) is allocated in a proportional manner across its neighbors in the network induced by the liability matrix. A pro-rata allocation is a standard allocation mechanism in financial debt theory [Eisenberg and Noe 2001; Fabozzi and Markowitz 2002; Rogers and Veraart 2013; Glasserman and Young 2016], and we adopt this standard in our article while allocating service liability debts. Given a solution to (2), an organization i is said to be *dysfunctional* if $p_i(\vec{x}) < \bar{p}_i(\vec{x})$ implying that its assets

are less than the liability it owes to other organizations in the service chain network. We illustrate the basic working of the GSOM model via an example in the Appendix.

3.3 Uniqueness of the Clearing Vector

Here, we investigate on the uniqueness of clearing vector obtained as the solution to the fixed point equation in (2). In this regard, we have the following theorem.

Theorem 3.1. *The clearing vector is unique if from every organizational node i there exists a chain of positive obligations to some organizational node k that has positive obligations to itself.*

Proof. It follows from [Eisenberg and Noe 2001] that a solution to (2) can be constructed iteratively as follows. Given a vector $\vec{x}$, define the mapping $\Phi : \mathrm{R}_+^n \to \mathrm{R}_+^n$ as

$$\Phi_i(\vec{p}) = \bar{p}_i \wedge \left(\sum_j p_j a_{ji} + c_i - x_i \right)_+ . \tag{3}$$

Starting with $\vec{p^0} = \bar{p}$, let $\vec{p^1} = \Phi(\vec{p^0})$, $\vec{p^2} = \Phi(\vec{p^1}), \ldots$, and so on. This iteration yields a monotone decreasing sequence $\vec{p^0} \geq \vec{p^1} \geq \ldots\ldots$. Since it is bounded below, it has a limit p', and since Φ is continuous p' satisfies (2). Hence it is a clearing vector. We now claim that p' is in fact the only solution to (2). Suppose by way of contradiction that there is another clearing vector, say $p'' \neq p'$. Then, the net worth of all organizational nodes must be the same under the two vectors, i.e.,

$$p'A + (c - x) - p' = p''A + (c - x) - p''.$$

Rearranging the terms, it follows that

$$(p'' - p')A = p'' - p'; \; p'' - p \neq 0.$$

This means that the matrix A has Eigenvalue 1, which is impossible because under our assumption A has spectral radius less than 1 - equivalent to the condition that, from every organizational node i, there exists a chain of positive obligations to some organizational node k that has positive obligations to itself. Thus, p' is the only solution to (2) and equivalently the clearing vector is unique. □

Theorem Implication. The uniqueness of the clearing vector provides the benefit of practically dealing with a single vector of liability payments, over the challenge of computationally searching for multiple vectors. The assumption that matrix A has a spectral radius less than 1 is quite practical in the sense that a chain of obligations ending in an obligation loop around the same organization, i.e., self-liability, is common in practice. As an example, the concept of self-liability could arise in the context of the popular notion of self-insuring an organization, which is common in business sectors.

4 ESTIMATING BLACKOUT CHANCE - SINGLE SOURCE CASE

In this section, we estimate the impact of the underlying service network topology the cyber-loss distribution on the probability of a cascading cyber-blackout among a given set of organization nodes when a single source node becomes dysfunctional to provide service in the aftermath of a cyber-attack. This section is divided into three main parts: in the first part, we provide a non-trivial general estimate of network impact on cyber-blackout probability irrespective of the loss distribution function; in the second part, we estimate the network impact on cyber-blackout probability for a certain popular family of proportional (as a function of node revenue base) loss distributions, i.e., the Beta distribution; finally, we study the effect of the network impact on cascading

Table 1. Table of Important Notations

N	set of organizations
$\bar{P} = (\bar{p}_{ij})$	payment matrix
$\overrightarrow{c}$	vector of wealth held by each organization i
$\overrightarrow{b}$	vector of liability-free losses
$\bar{P}_i$	net liability of i
X_i	random variable representing loss to i on random shock
$F(x_1, \ldots . x_n)$	joint cdf of losses at organization
$A = (a_{ij})$	relative liability matrix
$\overrightarrow{P}(\overrightarrow{x})$	clearing vector
β_i	proportion of i's monetary liabilities to other nodes
λ_i	leverage ratio of organization i w.r.t. c_i
$\overrightarrow{S}$	vector of monetary shortages at organizations

cyber-blackout probability for various non-heavy tailed distributions. A collection of important notations used in the article is provided in Table 1.

4.1 Analysis Setup

We first re-iterate the definition of the terms *organizational dysfunction* and *cyber-blackout*. As previously mentioned, organizational dysfunction happens when a given organization is unable to provide service to customers who rely on the former. Service could be of myriad forms, one popular example being the ability to protect customer private information; another example of a critical nature being the ability to provide non-interrupted energy to different customer segments in the industry. Potential impact of organizational dysfunction could result in monetary business losses, and loss of reputation resulting in loss of business. A cyber-blackout happens when individual organization dysfunction contributes to a cascade *(due to a contagion effect)* of organizational dysfunctions, where each subsequent organization that became dysfunctional was relying on other organizations that had already become dysfunctional. Note here that an organization could be a single user as well. A practical example of a cyber-blackout is service disruption in a power grid caused by a cyber-attack which in turn causes a cascade of power unavailability issues in different sectors (e.g., manufacturing, transportation) of the industry, thereby leading to business disruptions that cause commercial losses. *In our work, we characterize dysfunctionality in a monetary fashion by mapping it to the case when the monetary value of the available resources of an organization is less than what it owes other organizations (in the event of their unability to provide service) which are liable on the former for service.* More formally, given a solution to (2), an organization i is said to be *dysfunctional* if $p_i(\vec{x}) < \bar{p}_i(\vec{x})$ implying that its assets are less than the liability it owes to other organizations in the service chain network.

In order to formulate our results, we need the following notation. Let D be the set of nodes that we are interested in investigating whether they can go dysfunctional due to a cascading disruption effect resulting from a cyber-attack on a given source node i that made i dysfunctional. Let $\beta_i = \frac{\bar{p}_i}{b_i + \bar{p}_i}$ be the proportion of node i's service-related monetary liabilities to other organizational entities (nodes) in the system. We assume that $\beta_i > 0$, i.e., each node has a non-zero service liability external to itself. Recall that $w_i > 0$ is node i's initial net worth in the aftermath of it being subject to a cyber-incident, and c_i represents the vector of wealth/resources held by node i that is *not yet* subject to cyber-attack. *We assume that $w_i < c_i$, since otherwise i could never go dysfunctional*

directly through losses in c_i post a cyber-attack that affects c_i. We define the ratio $\lambda_i = \frac{c_i}{w_i} \geq 1$ to be the leverage ratio of i with respect to c_i, and denotes the vulnerability of i - more the c_i. In practice, λ_i denotes the dependency of i on other nodes (organizations) in the network to maintain its QoS with clients, and consequently serves as a vulnerability index of i. Thus, the greater the value of λ_i (implying lower net worth), the higher the chances of it becoming dysfunctional after other nodes in the system are cyber-attacked. A useful example to think of in this regard is a business that rents virtual machines from a third-party cloud provider to generate a significant fraction of its revenue, and the former gets dysfunctional via a DDoS attack. As a result, the net worth of the cloud-driven business organization falls significantly, thereby increasing its leverage index.

4.2 Estimate of Blackout Probability

In this section, given D and a node $i \notin D$, we first derive a general estimate of the cyber-blackout probability without taking into account specific forms of loss distribution functions. In this regard, we are interested in two quantities: *(i) a probability estimate that all organizations in D go dysfunctional, and (ii) the mathematical condition which guarantees the impossibility of a cascading effect from i to D*. The first quantity has implications to a cyber-insurer in the insurance industry who might be responsible for covering aggregate or accumulative risks of the organizations in D, and the value of this quantity will help the insurer design and manage its portfolio mechanisms to prevent it from going bankrupt. The second quantity has implications on individual organizations regarding boosting their investments in cyber-security so much as to prevent them getting dysfunctional and subsequently saving face and money, and furthermore arresting a cascading service disruption process.

We have the following proposition regarding a general *bound-based* estimate of the cyber-blackout probability independent of the specific forms of loss distribution functions.

Proposition 4.1. *Suppose that only organizational node i suffers a loss in its c_i from a cyber-attack, i.e., $x_j = 0$, $\forall j \neq i$, and that no organization is dysfunctional prior to i suffering the loss. Fix a set of nodes D not containing i. The probability that the loss causes all nodes in D to become dysfunctional is* ***upper bounded*** *by*

$$P\left(X_i \geq w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j\right). \tag{4}$$

A cascading effect from i to D is impossible if

$$\sum_{j\in D}\frac{w_j}{w_i} > \beta_i(\lambda_i - 1). \tag{5}$$

Proof. Let $D(\vec{x}) \equiv \bar{D}$ be the dysfunctional set resulting from the loss vector X, whose coordinates are all zero except X_i. By assumption, i causes other nodes to become dysfunctional; hence, i itself must become dysfunctional, i.e., $i \in \bar{D}$. To prove (4), it suffices to show that

$$\beta_i(X_i - w_i) \geq \sum_{j\in\bar{D}-\{i\}} w_j \geq \sum_{j\in D} w_j. \tag{6}$$

The second inequality in (6) follows from the assumption that no nodes are in default before the loss and the fact that we must have $d \subseteq \bar{D} - \{i\}$ for all nodes in D to default. For the first inequality in (6), define the *shortage* at organizational node j to be the difference $s_j = \bar{p}_j - p_j$. From (2), we see that the vector of shortages $\vec{s}$ satisfies

$$\vec{s} = (\vec{s}A - w + X)_+ \wedge \bar{p}.$$

Using (3), we have $s_j > 0$ for $j \in \bar{D}$ and $s_j = 0$ otherwise. We use a subscript $\bar{D}$ as in $s_{\bar{D}}$ or $A_{\bar{D}}$ to restrict a vector or matrix to the entries corresponding to nodes in the set $\bar{D}$. Then, the vector of shortages as the nodes of $\bar{D}$ satisfies

$$s_{\bar{D}} \leq s_{\bar{D}} A_{\bar{D}} - w_{\bar{D}} + X_{\bar{D}}, \tag{7}$$

hence,

$$X_{\bar{D}} - w_{\bar{D}} \geq s_{\bar{D}}(I_{\bar{D}} - A_{\bar{D}}). \tag{8}$$

The vector $s_{\bar{D}}$ is strictly positive in every coordinate. From the definition of β_j, we also know that the jth row sum of $I_{\bar{D}} - A_{\bar{D}}$ is at least $1 - \beta_j$. Hence,

$$s_{\bar{D}}(I_{\bar{D}} - A_{\bar{D}}) \cdot 1_{\bar{D}} \geq \sum_{j \in \bar{D}} s_j(1 - \beta_j) \geq s_i(1 - \beta_i). \tag{9}$$

From (7), it follows that the shortage at node i is at least as large as the initial amount by which i becomes dysfunctional, that is,

$$s_i \geq X_i - w_i > 0. \tag{10}$$

From (8)–(10), we can conclude that

$$\sum_{j \in \bar{D}} (X_j - w_j) \geq s_i(1 - \beta_i) \geq (X_i - w_i)(1 - \beta_i). \tag{11}$$

This establishes (6) and the first statement of the proposition. The second statement follows from the first by recalling that the loss to c_i's cannot exceed their value, i.e., $X_i < c_i$. Therefore, by (4), the probability of contagion in the context of organizational dysfunctionality is zero if

$$c_i \leq w_i + \frac{1}{\beta_i} \sum_{j \in D} w_j.$$

Dividing through by w_i, we see that this is equivalent to the condition

$$\sum_{j \in D} \frac{w_j}{w_i} > \beta_i(\lambda_i - 1),$$

which is the second statement of the proposition. Hence, we have proved Proposition 4.1. □

Proposition Implication. Note that the bounds in the theorem are completely general and do not depend on the distribution of the losses, or on the network topology. The condition in (5) is intuitive and states that dysfunction contagion from i to D is impossible if the total net worth of the nodes in D is sufficiently large (could be made possible by making proper investments in cyber-security) relative to the net worth of i weighted by (a) the exposure of the system to organizational node i as measured by β_i and (b) the vulnerability of i as measured by the leverage ratio λ_i. More generally contagion will be weak if unless originating node is highly leveraged and has a relative high proportion of obligations to other nodes (e.g., if originating node is an organization like Amazon providing cloud services to multiple other organizations [Wootton 2017; AWS Sales 2018]). A similar interpretation applies to (4).

Cyber-Insurance Perspective. In the context of cyber-insurance, the proposition implies that insurers should incentivize organizations (through appropriate contract design) to boost up their cyber-hygiene so that an organization's net worth (the denominator of the leverage term) is high. This implies that even if individual node dependencies are high, each is secure enough to a threshold that significantly dampens the chance of its revenue base to be shrink via a cyber-attack. The incentive problem has been a challenging one in the cyber-insurance space, and one particular solution direction for networks has recently been explored by Pal et al. [2017].

Contagion vs Independent Losses. We now investigate results tying the probability of a cyber-blackout through a contagion from a given organizational node i to a given subset D of nodes, to the probability of the same under direct independent losses (e.g., losses incurred by organizations due to cyber-attacks that take advantage of poor cyber-hygiene practices in the organizations) experienced at the nodes. *We say that the contagion effect with respect to organizational node dysfunction is weak if*

$$P\left(X_i \geq w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j\right) \leq P(X_i > w_i)\prod_{j\in D} P(X_j > w_j). \tag{12}$$

The expression on the left bounds the probability that nodes in D become dysfunctional through contagion from i, while the expression on the right is the probability (computed using the loss distribution for individual nodes) that the same nodes become dysfunctional through independent direct losses. The intuition for weak contagion is as follows: the RHS of the expression has the product of events, which means we consider the case where all nodes in D and node i get dysfunctional in an independent fashion. Thus, we have a sequence of 'less than 1' terms making the RHS smaller and smaller, yet it never gets small enough to become smaller than the LHS, that represents the network contagion effect. It goes without saying that the inequality depends heavily on w_i and β_i, and specific conditions in this regard are stated in the implications of Theorem 4.2 (see later). Note that in practice the assumption of direct independent losses is somewhat unrealistic: in practice one would expect the losses to different nodes be positively associated (correlated). In that case, we observe *a fortiori* that the probability of organizational dysfunction is weak even if the cyber-losses are positively correlated, and the above equation would hold here as well. That is, $P(X_i > w_i)\prod_{j\in D} P(X_j > w_j) < P(X_i > w_i, X_j > w_j, \forall j \in D)$. *We say that the contagion effect is strong if*

$$P\left(X_i \geq w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j\right) > P(X_i > w_i)\Pi_{j\in D} P(X_j > w_j). \tag{13}$$

The intuition for a strong contagion effect is just the converse of that for a weak contagion effect.

4.3 Distribution-Based Estimate of Cyber-Blackout Probability

Having provided a general estimate of the cyber-blackout probability, we now estimate this probability under the effect of a given loss distribution across different nodes in the organizational network. Let us assume that the cyber-losses at a given organizational node i scales with the portfolio c_i of the organization. Based on recent data from Symantec [Symantec 2016], this is a reasonable assumption to make irrespective of whether cyber-attackers target organizations big or small. Let us also assume that the distribution of these *relative losses*, i.e., with respect to c_i, is the same for the nodes, and *independent* among the nodes (note that this *does not* imply absolute losses are independent).

Then, there exists a distribution function $H : [0, 1] \rightarrow [0, 1]$ such that

$$F(x_1, \ldots\ldots, x_n) = \Pi_{1\leq i\leq n} H\left(\frac{x_i}{c_i}\right). \tag{14}$$

Beta distributions provide a flexible standard family with which to model the distribution of relative losses that lie in the interval $[0, 1]$, and generalizes other distributions that work with bounded

intervals [Johnson et al. 1995]. We work with *Beta densities* of the form

$$h_{p,q} = \frac{y^{p-1}(1-y)^{q-1}}{B(p,q)},\ 0 \le y \le 1,\ p,q \ge 1, \tag{15}$$

where $B(p,q)$ is a normalizing constant. Note that (15) is general enough to allow a mode anywhere in the unit interval. The subset with $p = 1$, $q > 1$ has a decreasing density and seems the most realistic, whereas the subset with $q = 1, p > 1$ has an increasing density and could be considered "heavy-tailed" in the sense that it assigns greater probability to greater losses. We have the following result regarding a distribution specific estimate of the cyber-blackout probability.

Theorem 4.2. *Assume relative loss distributions across all organizational nodes are i.i.d. Beta distributed, and the net worth of every node is initially non-negative. Let D be a non-empty subset of nodes and let $i \notin D$. Then, a contagion effect with respect to organizational dysfunction is impossible if*

$$\sum_{j\in D} w_j > w_i\beta_i(\lambda_i - 1), \tag{16}$$

and is weak if

$$\sum_{j\in D} w_j \ge w_i\beta_i \sum_{j\in D} \frac{\lambda_i - 1}{\lambda_j}. \tag{17}$$

Proof. Proposition 4.1 implies that contagion is weak from i to D if

$$P\left(X_i \ge w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j\right) \le P(X_i > w_i)\Pi_{j\in D}P(X_j > w_j). \tag{18}$$

On the one hand, this certainly holds if $w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j > c_i$, for then contagion is impossible. In this case, we obtain, as in (5)

$$\sum_{j\in D} \frac{w_j}{w_i} > \beta_i(\lambda_i - 1). \tag{19}$$

Suppose, on the other hand, that $\left(w_i + \frac{1}{\beta_i}\sum_{j\in D} w_j\right) \le c_i$. By assumption, the relative losses $\frac{X_k}{c_k}$ are independent and beta distributed as in (15). In the uniform case $p = q = 1$, Equation (18) is equivalent to

$$\left[1 - \left(\frac{w_i}{c_i} + \frac{1}{\beta_i c_i}\right)\sum_{j\in D} w_j\right] \le \left(1 - \frac{w_i}{c_i}\right)\Pi_{j\in D}\left(1 - \frac{w_j}{c_j}\right). \tag{20}$$

We claim that (20) implies (18) for the full family of Beta distributions in (16). To see why, first observe that the cumulative distribution $H_{p,q}$ of $h_{p,q}$ satisfies

$$1 - H_{p,q}(y) = H_{q,p}(1-y).$$

Hence, (18) holds if

$$H_{q,p}\left(1 - \frac{w_i}{c_i} - \frac{1}{\beta_i c_i}\sum_{j\in D} w_j\right) \le H_{q,p}\left(1 - \frac{w_i}{c_i}\right)\Pi_{j\in D}H_{q,p}\left(1 - \frac{w_j}{c_j}\right). \tag{21}$$

But (21) follows from (20) because Beta distributions with $p, q \ge 1$ have the submultiplicative property

$$H_{q,p}(xy) \le H_{q,p}(x)H_{q,p}(y),\ x, y \in [0,1].$$

It therefore suffices to establish (21), which is equivalent to

$$\frac{1}{\beta_i c_i}\sum_{j\in D} w_j \geq \left(1-\frac{w_i}{c_i}\right)\left(1-\Pi_{j\in D}\left(1-\frac{w_j}{c_j}\right)\right). \tag{22}$$

Given any real number $\theta_j \in [0,1]$, we have the inequality

$$\Pi_j(1-\theta_j) \geq 1-\sum_j \theta_j. \tag{23}$$

Hence, a sufficient condition for (22) to hold is that

$$\frac{1}{\beta_i c_i}\sum_{j\in D} w_j \geq \left(1-\frac{w_i}{c_i}\right)\sum_{j\in D}\frac{w_j}{c_j}. \tag{24}$$

After rearranging the terms and using the fact that $\lambda_k = \frac{c_k}{w_k}$ for all k, we obtain (17). This concludes the proof of Theorem 4.2. □

From the argument in (21), it is evident that the same result holds if the losses to each node j are distributed with parameters p_j, q_j in (15) with $p_i \leq \min_{j\in D} p_j$ and $q_i \geq \max_{j\in D} q_j$.

Theorem Intuition. As noted in Proposition 4.1, the condition in (16) states that the contagion from i to D is impossible if the total net worth of the nodes in D is sufficiently large (could be made possible by making proper investments in cyber-security) relative to the net worth of i weighted by (a) the exposure of the system to organizational node i as measured by β_i, and (b) the vulnerability of i as measured by the leverage ratio λ_i. The condition in (17) compares the total net worth of D relative to that of i with the leverage ratio of i relative to that of the nodes in D. With other parameters held constant, increasing the relative net worth of D (again via making higher investments in security) makes contagion weaker in the sense that it strengthens the inequality; increasing the leverage ratio of i relative to that of the nodes in D has the opposite effect because there is higher potential (and impact) to target unattacked resources worth c_i and also those whose net worth is low. Importantly, the two effects are mediated by β_i —a lower β_i makes D vulnerable to i and makes D less sensitive to the degree of leverage at i. Now recalling that $\lambda_j = \frac{c_j}{w_j}$, we can write (17) in the following equivalent form:

$$\frac{\sum_{j\in D} c_j \lambda_j^{-1}}{\sum_{j\in D}\lambda_j^{-1}} \geq c_i\beta_i(1-\lambda_i^{-1}). \tag{25}$$

Written this way, the condition states that contagion from i to D is weak if the average size of the nodes in D weighted by their inverse leverage ratios is sufficiently large relative to i —evident as a result of high net worth of nodes in D. On the right side of the inequality in (25), $c_i\beta_i$ measures the organizational system's exposure to node i's assets worth c_i, and the factor $(1-\lambda_i^{-1})$ is greater when node i is more highly leveraged. Thus, inequality (25) is harder to satisfy, and D is more vulnerable to contagion from i, if large (high asset) nodes in D are more highly leveraged, or if node i is more highly leveraged.

Theorem Implications. A key implication to Theorem 4.2 is that without substantial node heterogeneity (see Corollary 4.4 for specific mathematical conditions), contagion with respect to organizational dysfunction will be weak, irrespective of the structure of the network induced by the liability matrix (also validated experimentally on real and synthetic data in Section 6). More generally, from Proposition 2, contagion will be weak unless the originating node is highly leveraged and has a relative high proportion of obligations to other nodes. Consequently, with respect to node heterogeneity, the following result is immediately obvious.

Corollary 4.3. *Assume that all nodes i have the same value c for c_i. Under the assumptions of Theorem 3.1, contagion is weak from any node to any other set of nodes.*

Proof. The result follows from the fact that $\beta_i(1-\lambda_i^{-1}) < 1$, and the fact that when $c_i = c$, for all i, (25) holds for all i and D. Thus, we have proved Corollary 4.3. □

The implication of this corollary is that organizational heterogeneity with respect to resources characterized by c_i's is a necessary condition (though not sufficient) for a cascading service disruption effect to take place. *Since, in reality, organizations are heterogeneous, cyber-blackouts are possible, though under certain conditions (see Corollary 4.4).*

Now suppose that $c_1 \geq c_2 \geq \cdots \geq c_m$, Since losses are proportional to c_i, a loss to c_1 maximizes the contagion to other nodes. This fact is formalized via the following corollary:

Corollary 4.4. *If $c_1 \geq c_2 \geq \cdots \geq c_m$, then contagion from organizational node 1 to nodes 2,.....,m is weak if $c_2 \geq \beta_1(c_1 - w_1)$ and $c_j \geq (c_{j-1} - w_{j-1})$, $j = 2, \ldots, m$, strong otherwise. Contagion is impossible if $c_2 - c_m + w_m > \beta_1(c_1 - w_1)$.*

Proof. The result directly follows from (20) and (22). □

The implications of this corollary are that the lower bounds for c_j ensure that the potential spillovers from other nodes cannot lead to the full set D of nodes into dysfunction *regardless of the liability network topology.* This does not imply that the network structure has no effect on the probability of contagion—it just showcases the fact that in quite a few situations the probability of contagion with respect to organizational dysfunction will be lower than the probability of an organization being rendered dysfunctional due to direct losses (see Section 6 for an experimental validation).

Cyber-Insurance Perspectives. In the context of cyber-insurance, the implications of Proposition 2 carry over, in addition to Theorem 4.2 and the subsequent corollaries bolstering the future increase in global cyber-insurance market valuation, as risk-averse insurers would not have to worry much about strong contagion effects in selling cyber-insurance policies. In addition, the common knowledge among organizations about inevitable node heterogeneity with respect to monetary assets, will psychologically lead them to invest in cyber-insurance as well as security enhancing practices due to a certain fear of risk cascading.

With respect to the scale of aggregate risk coverage burden on an insurance company, cyber-blackouts may be quite unlikely if set D is large, which reduces the likelihood of an insurance agency going bankrupt, and this implication holds irrespective of the underlying liability network topology. When D is a small set of heterogeneous organizations, an insurance company is also less likely to be bankrupt, even if some organizations in the set are large-sized and incur large losses. Now as for the case of simultaneous independent direct losses on *all* of the organizations in D which might positively contribute to cyber-blackout probability—in practice, this is a very low probability event for large-sized D.

4.4 Extending the Distribution Space

A drawback with the *Beta* distribution is that the probability of c_i going to zero in the aftermath of a cyber-attack, is zero. This clearly may not be true in practice and we cannot rule out the (potentially futuristic) scenario where the c_i's could be wiped out due to a big cyber-hit—something analogous to a cyber 9/11. In this section, we aim to extend our analysis by accounting for popular loss distributions other than the Beta distribution, that do not suffer from the above-mentioned drawback.

In order to capture the non-positive probability of c_i's going to zero, we propose the following model: Let $X_i^0 \geq 0$ be a *primary loss* (potentially unbounded in size) and let $X_i = (c_i \wedge X_i^0)$ be the

resulting loss to c_i for organizational node i —i.e., we truncate the loss to put mass at c_i thereby setting up the way to assign positive probability to c_i going to zero. Assume that the primary losses have a joint distribution of the form:

$$F^0(x_1^0, \ldots\ldots, x_n^0) = \Pi_{1\le i\le n} H^0\left(\frac{x_i^0}{c_i}\right), \tag{26}$$

where H^0 is a distribution function on the non-negative real line. More specifically, we assume (for now) that the primary losses are i.i.d. and that a given x_i^0 affects every unit of c_i equally. A random variable with distribution function G and density g is said to have an increasing failure rate (IFR) distribution if $\frac{g(x)}{1-G(x)}$ is an increasing function of x. Given the assumption that $Y_i = \frac{X_i^0}{c_i}$ are i.i.d., Y_i's are IFRs. Other popular examples of IFR's are *normal*, *exponential*, and *uniform distributions*; more generally, all *log-concave* distributions. Our model showcases the IFR property common to multiple popular distribution families and helps us extend results in the previous section to distributions beyond the Beta distribution. We have the following result regarding a non-specific distributional, i.e., IFR-distributed estimate of the cyber-blackout probability.

Theorem 4.5. *Assume relative primary loss distributions across all organizational nodes are i.i.d. IFR-distributed, and the net worth of every node is initially non-negative. Let D be a non-empty subset of nodes and let $i \notin D$. Then, a contagion effect with respect to organizational dysfunction is impossible if*

$$\sum_{j\in D} w_j > w_i\beta_i(\lambda_i - 1), \tag{27}$$

and is weak if

$$\sum_{j\in D} w_j \ge w_i\beta_i \sum_{j\in D}\frac{\lambda_i}{\lambda_j}. \tag{28}$$

Proof. Through relabeling, we can assume that the source node for contagion is $i = 1$ and that the infected nodes are $D = \{2, 3, \ldots, m\}$. By Proposition 1, we know that contagion is weak from 1 to D if

$$P\left(X_1 > w_1 + \frac{1}{\beta_1}\sum_{2\le j\le m} w_j\right) \le \Pi_{1\le j\le m} P(X_j > w_j). \tag{29}$$

Since $X_1 = c_1 \wedge X_1^0$, the left-handed side is zero when $w_1 + \frac{1}{\beta_1}\sum_{2\le j\le m} w_j > c_1$. Thus, contagion is impossible if

$$\sum_{2\le j\le m}\frac{w_j}{w_1} > \beta_1(\lambda_1 - 1). \tag{30}$$

Let us therefore assume that $w_1 + \frac{1}{\beta_1}\sum_{2\le j\le m} w_j \le c_1$. Define the random variables $Y_i = \frac{X_i^0}{c_i}$. Then, the weak contagion from 1 to D holds if

$$P\left(Y_1 > \frac{w_1}{c_1} + \frac{1}{\beta_1 c_1}\sum_{2\le j\le m} w_j\right) \le \Pi_{1\le j\le m} P\left(X_j > \frac{w_j}{c_j}\right), \tag{31}$$

where the latter holds from the assumption that Y_i are i.i.d. By assumption that Y_1 is IFR, hence we have from Barlow and Proschan [1975]

$$P(Y_1 > s + t | Y_1 > s) \le P(Y_1 > t),\ \forall s, t \ge 0.$$

It follows that

$$P\left(Y_1 > \sum_{1\le k\le m} \frac{w_k}{c_k}\right) \le \Pi_{1\le j\le m} P\left(X_j > \frac{w_j}{c_j}\right). \tag{32}$$

Together with (31), it shows that contagion from 1 to D is weak provided that

$$P\left(Y_1 > \frac{w_1}{c_1} + \frac{1}{\beta_1 c_1} \sum_{2\le j\le m} w_j\right) \le P\left(Y_1 > \sum_{1\le k\le m} \frac{w_k}{c_k}\right). \tag{33}$$

This clearly holds if

$$\frac{w_1}{c_1} + \frac{1}{\beta_1 c_1} \sum_{2\le j\le m} w_j \ge \sum_{1\le k\le m} \frac{w_k}{c_k}, \tag{34}$$

which is equivalent to

$$\frac{1}{\beta c_1} \sum_{2\le j\le m} w_j \ge \sum_{2\le j\le m} \frac{w_j}{c_j} = \sum_{2\le j\le m} \lambda_j^{-1}. \tag{35}$$

Since $c_1 = \lambda_1 w_1$, we can rewrite (35) as

$$\sum_{2\le j\le m} \frac{w_j}{w_1} \ge \beta\lambda_1 \sum_{2\le j\le m} \lambda_j^{-1}. \tag{36}$$

We have therefore shown that if contagion from 1 to $D = \{2, 3, \ldots, m\}$ is possible at all, then Equation (36) is a sufficient condition for weak contagion with respect to service dysfunctionality. From (36), we see that a simple sufficient condition for weak contagion is $c_j \ge \beta_1 c_1$, $j = 2, \ldots, m$, and the condition $\sum_{j=2}^{m} w_j > \beta_1(c_1 - w_1)$ make contagion impossible. Thus, we have proved Theorem 4.5. □

Theorem Intuition and Implications. We have the following very powerful system implication as a result of Theorem 4.5 and Corollary 4.4, given a single node i (that got hit by a cyber-attack) and an organization set D —*the conditions for weak contagion and the impossibility of contagion with respect to organizational service dysfunction is the same irrespective of the loss distributions and the underlying network topology, as long as the distributions satisfy the general IFR property.* Thus, in a sense, the specificity of loss distributions is "irrelevant" to the conditions necessary for a cyber-blackout. The intuition is similar to that of Theorem 4.2. With respect to node heterogeneity, the following result is immediately obvious from Theorem 4.5.

Corollary 4.6. *Assume that all nodes i have the same value c for c_i. Under the assumptions of Theorem 4.2, contagion is weak from any node to any other set of nodes.*

Proof. It is evident upon writing (36) as

$$\frac{\sum_{j\in D} c_j \lambda_j^{-1}}{\sum_{j\in D} \lambda_j^{-1}} \ge \beta_i c_i.$$

Hence, we have proved Corollary 4.6. □

Cyber-Insurance Perspective. With respect to cyber-insurance, the implications of Theorem 4.5 are the same as those from Theorem 4.2.

Non-i.i.d Primary Losses. In the beginning of this section, we had assumed that primary losses across organizational nodes are i.i.d. However, this assumption is conservative in practice. Here, we provide the conditions for weak contagion for specific but practical loss variables characterized by a Pareto-like or a heavy-tailed densities of the form $P(X_i > x) \approx ax^{-\mu}$ for some positive constants a and μ. First, we generate dependent random variables from independent random variables via

a standard statistical procedure as follows: let $Y_1, \ldots, Y_m$ be independent random variables, each distributed as t_ν —the Student t distribution with $\nu > 2$ degrees of freedom. Let $\hat{Y}_1, \ldots \hat{Y_m}$ have a standard multivariate Student t distribution with t_ν marginals. Clearly, $\hat{Y}_j$'s are uncorrelated but not independent. In order to make losses positive, we set $\tilde{X_j} = \tilde{Y_j}^2$, for each j, where $\tilde{X_j}$ has a Pareto-like tail.

Proposition 4.7. *With dependent primary losses, $\tilde{X_i}$,*

$$P\left(\tilde{X_i} > \sum_{j=1}^{m} w_j\right) \leq P(\tilde{X_j} > w_j,\ j = 1, \ldots, m),$$

for all $w_j \geq 0$, $j = 1, 2, \ldots, m$.

Proof. The proof follows via a direct application of Bound II for the F distribution (see Marshall et al. [1974]). □

The proposition implies that even with heavy-tailed losses, we may find that service dysfunction of a set of nodes through contagion originating from a single organizational node is less likely than service dysfunction via direct losses to individual nodes, if the losses are dependent.

5 EXPANDING ATTACK SOURCE AND TARGET SETS

In the previous section, we studied the impact of the dysfunctionality of a single organizational node on another *target set D* of organizational nodes. In this section, we study the impact of (multiple) successful cyber-attacks on the *entire organizational network*. More specifically, we model our goal as an estimate of the effect of the underlying liability network on the *net losses* in the overall system due to (simultaneous) successful cyber-attacks on c_i's of different organizational nodes i. In this regard, we first need to form a measure of the *total systemic impact* of loss due to cyber-attack. In this work, we shall take the systemic impact of a loss to be the total loss in value summed over all organizational nodes in the network. Given a loss realization $\vec{x}$, the total reduction in resources (assets) across all nodes in the network is

$$\sum_i x_i + S(\vec{x});\ S(\vec{x}) = \sum_i (\bar{p}_i - p_i(\vec{x})). \tag{37}$$

The term $\sum_i x_i$ is the direct loss in value from reductions in liability payments to the i's from their external network environment. The term $S(\vec{x})$ is the indirect loss in value from reductions in liability payments by the nodes to other nodes as well to themselves (due to self-liability). An overall measure of the riskiness of the network system is the expected loss in value, L, given by

$$L = \int \left(\sum_i x_i + S(\vec{x})\right) dF(\vec{x}). \tag{38}$$

The question we wish to examine is what proportion of these losses can be attributed to network connections between organizations?

5.1 Examination Setup

Let $\vec{x}$ be a loss value (instance) due to a cyber-attack, and correspondingly let $D = D(\vec{x})$ be the set of nodes that goes dysfunctional given $\vec{x}$. Under our assumptions, this set is unique because the clearing vector is unique. For notational simplicity, we suppress $\vec{x}$ in the ensuing discussion. As in the proof of Proposition 4.1, define the shortage in liability payments at organizational node i to

be $s_i = \bar{p}_i - p_i$, where $\vec{p}$ is the clearing vector. By definition of D, we have

$$s_i = \begin{cases} > 0, \ \forall i \in D \\ = 0, \ \forall i \notin D \end{cases}$$

Also, as in part of Proposition 4.1, let A_D be the $|D| \times |D|$ matrix obtained by restricting the relative liabilities matrix A to D and let I_D be the $|D| \times |D|$ identity matrix. Similarly, let $\vec{s}_D$ be the vector of shortages s_i corresponding to the nodes in D, let $\vec{w}_d$ be the corresponding net worth vector defined in (1), and let $\vec{x}_D$ be the corresponding vector of losses. The clearing condition in (3) implies the following equation, provided $s_i < \bar{p}_i$ (the condition that the net worth of any node is positive), for all i:

$$\vec{s}_D A_D - (\vec{w}_D - x_D) = \vec{s}_D. \tag{39}$$

Recall that A_D is substochastic, and by assumption, there exists a chain of obligations from any given node k to a node having strictly positive obligations to the itself. It follows that $\lim_{k\to\infty} A_D^k = 0_D$; hence, $I_D - A_D$ is invertible and

$$[I_D - A_D]^{-1} = I_D + A_D + A_D^2 + \ldots\ldots\ldots \tag{40}$$

From (39) and (40), we conclude that

$$\vec{s}_D = (\vec{x}_D - w_D)[I_D + A_D + A_D^2 + \ldots\ldots\ldots]. \tag{41}$$

Given a loss instance $\vec{x}$ with resulting dysfunctional organization set $D = D(\vec{x})$, define the vector $u(\vec{x}) \in \mathrm{R}_+^n$ such that

$$u_D(\vec{x}) = [I_D + A_D + A_D^2 + \ldots\ldots] \cdot 1_D, \ u_i(\vec{x}) = 0, \ \forall i \notin D. \tag{42}$$

Combining (37), (41), and (42) shows that the total losses for a given $\vec{x}$ can be expressed as

$$L(\vec{x}) = \sum_i (x_i \wedge w_i) + \sum_i (x_i - w_i) u_i(\vec{x}). \tag{43}$$

The first term represents the direct losses to remaining resources at each organizational node, and the second term represents the total shortage summed over all the nodes. The right side becomes an upper bound on $L(x)$ if $s_i = \bar{p}_i$ for some $i \in D(\vec{x})$. We call the coefficient $u_i = u_i(\vec{x})$ the *depth* of organizational node i in $D = D(x)$. The rationale for this terminology is as follows: Consider a Markov chain on D with transition matrix A_D. For each $i \in D$, u_i is the expected number of periods before exiting D, starting from node i. Expression (42) shows that *node depths measure the amplification of losses due to interconnections among nodes in the dysfunctional set.* We note here that the concept of node depth is dual to the notion of eigenvector centrality (or eigenvector-driven centrality measures) [Newman 2018]. To see the connection, let us restart the Markov chain uniformly in D whenever it exits D. This modified chain has an ergodic distribution proportional to $1_D \cdot [I_D + A_D + A_D^2 + \ldots\ldots]$ and its ergodic distribution measures the centrality of the nodes in D. It then follows that node depth with respect to A_D corresponds to centrality with respect to the transpose of A_D.

We can now bound the magnitude of the node depths in the dysfunctional set. We first define a set D of nodes to be *α-cohesive* if every node in D has at least α of its liabilities to other nodes in D, i.e., $\sum_{j\in D} a_{ij} \geq \alpha$, for every $i \in D$ [Morris 2000]. The *cohesiveness* of D is the maximum α, which we denote by α_D. As a lower bound for u_i, it follows from (42) that

$$u_i \geq \frac{1}{1 - \alpha_D}, \ \forall i \in D. \tag{44}$$

Thus, *the more cohesive the dysfunction set, the greater the depth of the nodes in that set and the greater the amplification of the associated loss.* We can also bound the node depths from above. Recall

that β_i is the proportion of i's liability to other nodes in the network. Let $\beta_D = \max\{\beta_i : i \in D\}$. We obtain the upper bound assuming $\beta_D < 1$ as follows:

$$u_i \leq \frac{1}{1-\beta_D}, \forall i \in D. \tag{45}$$

The bounds in (44) and (45) depend on the dysfunctional set D, which in turn depends on $\vec{x}$. A uniform upper bound is given by

$$u_i \leq \frac{1}{1-\beta^+}, \forall i \in D;\ \beta^+ = \max \beta_i < 1. \tag{46}$$

We are now in a position to compare the expected systemic losses in a given network of interconnections, with the expected losses without such interconnections, in order to gauge the effect of service disruption contagion in a network.

5.2 Comparing Expected Systemic Losses

Consider the following system setting as already discussed in Section 3. We fix a set of n organizational nodes, $N = \{1, 2, \ldots, N\}$, vectors $\vec{c}$, and $\vec{b}$ as before. Assume that the net worth w_i of node i is non-negative before a loss due to a cyber-attack is realized, and that liability network interconnections are determined via the $n \times n$ matrix $\bar{P}$. Let us now have another system setting where we eliminate all connections between nodes, i.e., let $\bar{P}^0$ be the $n \times n$ matrix of zeros. Each node i in this setting has resources c_i that are yet to be attack-targeted, and self-liabilities, b_i. In order to keep an organization's net worth unchanged, we introduce "fictitious" resources c_i and self-liabilities b_i to maintain balance. More specifically, if $c_i - b_i < w_i$, we give i a new class of resources in the resource amount $c_i' = w_i - (c_i - b_i)$. If $c_i - b_i > w_i$, we give i a new class of self-liabilities in the amount $b_i' = w_i - (c_i - b_i)$. We assume that the new resources are safe, i.e., they are not subject to cyber-attacks, and that the new liabilities have the same priority as other liabilities. Let $F(x_1, \ldots\ldots x_n)$ be a joint loss distribution that is homogeneous in resources, i.e., $F(x_1, \ldots\ldots, x_n) = G(\frac{x_1}{c_1}, \frac{x_2}{c_2}, \ldots\ldots \frac{x_n}{c_n})$, where G is a symmetric cumulative distribution function. We do not assume that losses across nodes are independent. We say that F is IFR if its marginal distributions are IFR; this is equivalent to saying that the marginals of G are IFR. Let $\bar{L}$ be the expected total losses in the original network and let $\bar{L}^0$ be the expected total losses when the connections are removed. We have the following result relating $\bar{L}$ and $\bar{L}^0$.

Theorem 5.1. *Let $N(\vec{b}, \vec{c}, \vec{w}, \bar{P})$ be an organizational network system and let N^0 be the analogous system with all the network connections removed. Assume that the loss distribution is homogeneous in resources and IFR. Let $\beta^+ = \max_i \beta_i < 1$, and let $\delta_i = P(X_i \geq w_i)$. Then, the ratio of expected losses in N^0 is at most*

$$\frac{\bar{L}}{\bar{L}_0} \leq 1 + \frac{\sum_i \delta_i c_i}{(1-\beta^+)\sum_i c_i}. \tag{47}$$

Proof. By assumption, the marginals of F are IFR distributed. A general property of IFR distributions is that "new is better than used in expectation", i.e.,

$$E[X_i - w_i | X_i \geq w_i] \leq E[X_i], \tag{48}$$

from Barlow and Proschan [1975]. It follows that

$$E[(X_i - w_i)^+] \leq P(X_i \geq w_i)E[X_i] = \delta_i E[X_i]. \tag{49}$$

By (43), we know that the total expected losses $\bar{L}$ can be bounded as

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + E\left[\sum_i (X_i - w_i)u_i(X)\right]. \tag{50}$$

From (46), we know that $u_i \leq \frac{1}{1-\beta^+}$ for all i; furthermore, we clearly have $X_i - w_i \leq (X_i - w_i)^+$ for all i. Thus,

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + (1-\beta^+)^{-1} \sum_i E[(X_i - w_i)^+]. \tag{51}$$

From this and (49), it follows that

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + (1-\beta^+)^{-1} \sum_i \delta_i E[X_i], \tag{52}$$

which reduces to

$$\bar{L} \leq \sum_i E[X_i] + (1-\beta^+)^{-1} \sum_i \delta_i E[X_i]. \tag{53}$$

When the network connections are excised, the expected loss is simply the expected sum of the losses, that is, $\bar{L^0} = \sum_i E[X_i]$. By the assumption of homogeneity in resources, we know that $E[X_i] \propto c_i]$ for all i. We conclude from this and (53) that

$$\frac{\bar{L}}{\bar{L_0}} \leq 1 + \frac{\sum_i \delta_i c_i}{(1-\beta^+)\sum_i c_i}.$$

Thus, we have proved Theorem 5.1. □

Theorem Implications. The theorem shows that increases in losses due to liability network interconnections will be very small unless β^+ (the maximum proportion of obligations by any node in the network) is close to 1, or the rate at which an organization becomes dysfunctional is high, both of which are quite unlikely in practice. Moreover, the latter statement also holds when the losses across nodes are dependent or correlated, regardless of the network structure.

Cyber-Insurance Perspective. Since the losses due to networked connectivity is primarily amplified due to a high β^+, which in turn implies the high dysfunctionality rate of an organization, it is imperative that cyber-insurers impose a strict control policy via their contracts with the organizations to ensure the highest standards of cyber-hygiene from the latter that results in low/moderate values of β^+. This in turn would reduce the probability of a cyber-blackout and also mitigate the chances of cyber-insurers going bankrupt in the process of covering correlated aggregate risk. An intuitively evident insurance policy mechanism in this regard is to premium discriminate between good hygiene and bad hygiene organizations [Pal et al. 2014]. Such policies have been shown to be market efficient in the economic sense.

5.3 Side-Effect (Reputational) Losses upon Organizational Node Dysfunctionality

In practice, once an organizational node becomes dysfunctional, as a side-effect, there is likely to be a negative impact on its reputation, apart from the usual loss in assets. Although such reputational losses are non-tangible in nature, we capture this phenomenon through a scalar multiplier $\rho \geq 0$ for each organization, reflecting on the amount of further reduction in its assets, on facing dysfunctionality. This amount is mathematically characterized as

$$\rho \left[\bar{p}_i - \left(c_i + \sum_{j \neq i} p_j a_{ji} - x_i \right) \right] \tag{54}$$

and reaches a maximum reduction amount where all the assets are wiped out. The term in the square is the difference between the dependency obligations of node i and its remaining assets. The parameter ρ magnifies the severity of the loss—the term in the square bracket, as a knock-on

effect due to a hit in organizational reputation. The resulting condition for a clearing vector is a solution to the following equation:

$$p_i(x) = \bar{p}_i \wedge \left[(1+\rho) \left(c_i + \sum_{j \neq i} p_j a_{ji} - x_i \right) - \rho \bar{p}_i \right]_+ \tag{55}$$

In terms of deficit $s_i = \bar{p}_i - p_i$, the above expression becomes

$$s_i = (1+\rho) \left[\sum_{j \neq i} s_j a_{ji} - w_i + x_i \right]_+ \wedge \bar{p}_i \tag{56}$$

We have the following result stating the effect of ρ on magnifying the deficit.

Theorem 5.2. *Let Φ^s_ρ denote the mapping from the vector s on the LHS of (56) to the s on the RHS of (56). Similarly, let Φ^p_ρ denote the mapping from the vector p on the LHS of (55) to the s on the RHS of (55). For any $\rho \geq 0$, the mapping Φ^s_ρ is monotone, increasing in ρ $\forall s$, bounded, and continuous on R^n_+, whereas the mapping Φ^p_ρ is monotone, decreasing in ρ $\forall p$, bounded, and continuous on R^n_+. Furthermore, Φ^s_ρ and Φ^p_ρ have least and greatest fixed points, s and p, respectively. The subsequent set of dysfunctional organizational nodes under minimal s and maximal p is increasing in ρ. In addition, the fixed point s and p is unique if $(1+\rho)A$ has a spectral radius < 1.*

Proof. The nature of mappings Φ^s_ρ and Φ^p_ρ follow directly from Theorem 1 of Eisenberg and Noe [2001]. Let $v_i = c_i + (pA)_i - x_i$. We have $\Phi^p_\rho(p)_i$ taking a value of $v_i - \rho(\bar{p}_i - v_i)$ when $v_i < \bar{p}_i$, and equalling $\bar{p}_i$ otherwise. This leads to the monotonicity of Φ^p_ρ with respect to ρ. The maximal fixed-point results as a limit of the sequence $\{\Phi^p_\rho\}$ beginning with $\bar{p}$, as a consequence of arguments in Section 3 of Eisenberg and Noe [2001]. More specifically, if $\rho_1 \leq \rho_2$, then $\Phi^p_{\rho_1} \leq \Phi^p_{\rho_2}$. The set of organizational nodes i for which $p_i < \bar{p}_i$ at the maximal fixed-point for ρ_1 must be contained within that for ρ_2. The uniqueness arguments follows from the case when there are no side-effect losses. The proof of statements regarding Φ^s_ρ follows exactly as those for Φ^p_ρ. Thus, we have proved Theorem 5.2. □

Theorem Implications. The obvious implication of the theorem is that reputation-centric losses on a cyber-attack increase contagion in a service network and results in an increased number of dysfunctional organizational nodes. The spectral condition for fixed-point uniqueness is a reality in the sense that a chain of obligations ending in an obligation loop around the same organizations, is common in practice.

Cyber-Insurance Perspectives. Reputation-centric losses are third-party in nature and only make it more difficult for profit-minded cyber-insurance agencies to cover amplified organizational losses. A step forward for the insurance companies to alleviate this concern is to charge amplified premiums based on the number of liabilities of a given organization.

Now suppose that the reputation-centric losses are not that big to wipe out all the assets of an organization. In that case, we have

$$S_D = (1+\rho)[S_D A_D - w_D + x_D].$$

Given, $I_D - (1+\rho)A_D$ is invertible, we have

$$s_D = (1+\rho)(x_D - w_D)[I_D - (1+\rho)A_D]^{-1}.$$

The deficit is given by

$$S(x) = S_D u_D = (1+\rho)(x_D - w_D)u_D(\rho), \tag{57}$$

where u_D is a modified organizational node depth vector given by $[I_D - (1+\rho)A_D]^{-1} \cdot 1_D$. When $(1+\rho)A_D$ has spectral radius 1, we get

$$u_D(\rho) = [I_D + (1+\rho)A_D + (1+\rho)^2 A_D^2 + \ldots..] \cdot 1_D.$$

Letting α_D denote the cohesiveness of D, the lower bound depth of each organizational node is given by $\frac{1}{1-(1+\rho)\alpha_D}$. Thus, from this expression of lower-bound depth, we can assess how side-effect losses deepen the total losses at dysfunctional organizational nodes. Let $\bar{L}$ be the expected total losses in the original network and let $\bar{L^0}$ be the expected total losses when the connections are removed. We have the following result relating $\bar{L}$ and $\bar{L^0}$.

Theorem 5.3. *Given* $(1+\rho)\beta^+ < 1$, *we have*

$$\frac{\bar{L}}{\bar{L_0}} \leq 1 + \frac{\sum_i \delta_i c_i}{(1-(1+\rho)\beta^+)\sum_i c_i}. \tag{58}$$

Proof. By assumption, the marginals of F are IFR distributed. A general property of IFR distributions is that "new is better than used in expectation", i.e.,

$$E[X_i - w_i | X_i \geq w_i] \leq E[X_i], \tag{59}$$

from Barlow and Proschan [1975]. It follows that

$$E[(X_i - w_i]^+] \leq P(X_i \geq w_i)E[X_i] = \delta_i E[X_i]. \tag{60}$$

By (43), we know that the total expected losses $\bar{L}$ can be bounded as

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + E\left[\sum_i (X_i - w_i)u_i(X)\right]. \tag{61}$$

From (46), we know that $u_i \leq \frac{1}{1-(1+\rho)\beta^+}$ for all i; furthermore, we clearly have $X_i - w_i \leq (X_i - w_i)^+$ for all i. Thus,

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + (1-(1+\rho)\beta^+)^{-1} \sum_i E[(X_i - w_i)^+]. \tag{62}$$

From this and (49), it follows that

$$\bar{L} \leq \sum_i E[X_i \wedge w_i] + (1-(1+\rho)\beta^+)^{-1} \sum_i \delta_i E[X_i], \tag{63}$$

which reduces to

$$\bar{L} \leq \sum_i E[X_i] + (1-(1+\rho)\beta^+)^{-1} \sum_i \delta_i E[X_i]. \tag{64}$$

When the network connections are excised, the expected loss is simply the expected sum of the losses, that is $\bar{L^0} = \sum_i E[X_i]$. By the assumption of homogeneity in resources, we know that $E[X_i] \propto c_i]$ for all i. We conclude from this and (53) that

$$\frac{\bar{L}}{\bar{L_0}} \leq 1 + \frac{\sum_i \delta_i c_i}{(1-(1+\rho)\beta^+)\sum_i c_i}.$$

Thus, we have proved Theorem 5.3. □

The theorem and cyber-insurance implications are similar to those of Theorem 5.1.

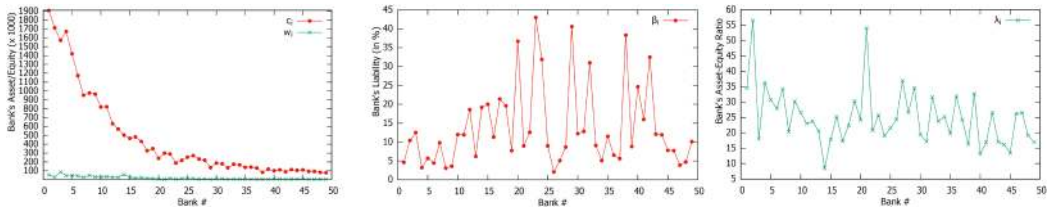

Fig. 1. Experimental parameters for real-world data: (a) c_i and w_i values (left), (b) β_i values (middle), and, (c) λ_i values (right).

6 EXPERIMENTAL EVALUATION

Experimenting with multiple real-world datasets related to cyber-attacks and their subsequent impact is an extremely difficult task, as data on cyber-security is really hard to obtain. As a result, in this section, we experiment on both real-world and synthetic banking sector application data obtained after a cyber-attack in the European Union. In this regard, we study the effects of parameters c_i, w_i, λ_i, and β_i on strong and weak contagion phenomena, in turn studying how our theoretical results apply in practice. We also experiment on synthetic data to study the effect of network topology on contagion phenomena.

6.1 Experimental Setup

One of the responsibilities of the European Banking Authority (EBA) is to ensure the orderly functioning and integrity of financial markets and the stability of the banking system in the EU. A primary supervisory tool to conduct such an analysis is via a stability test exercise. The aim of such a test is to assess the resilience of banking institutions to adverse market developments, as well as to contribute to the overall assessment of systemic risk in the EU banking system, where the systemic risk could be due to a cyber-attack. We collected data for a cyber-attack–induced stress test done in 2015. *Detailed information on inter-bank exposures needed to calibrate a full network was not publicly available. As a result, as aforementioned, we also generated 50 instances of synthetic random networks between banks in the 2015 dataset to study the effect of network topology on the contagion phenomena.*

For the real dataset, 90 banks from 21 countries participated in the stress test. For each bank, the EBA reports each bank's total exposure at the dysfunction state to other banks. The EAD measures a bank's total claims on all other banks, so we take this as the size of each bank's in-network assets. Subtracting this value from the total assets gives us c_i. For w_i (see Figure 1(a)), we use the equity values reported by the EBA, which then allows us to calculate $\lambda_i = \frac{c_i}{w_i}$ (see Figure 1(c)). The only remaining parameter we need is β_i, the fraction of a bank's liabilities owed to other banks. This information is not included in the EBA summary, nor is it consistently reported by the banks in their statements. As a rough indication, we assume that each bank's in-network liabilities equal its in-network assets.[2] This gives us $\beta_i = \frac{EAD}{\text{assets}-\text{equity}}$ (see Figure 1(b)). Some of the smallest banks have a problematic data, so as a simple rule we omit the ten smallest. We also omit any countries with only a single participating bank. This leaves us with 76 banks, out of which we work with 50 largest banks. For synthetic datasets, we estimate the parameters c_i, w_i, λ_i, and β_i in the same manner as for our real-world dataset.

[2]Based on Federal Reserve Release, the average value of β_i for commercial banks in the USA is about 3%, so our estimates for European banks would appear to be conservative.

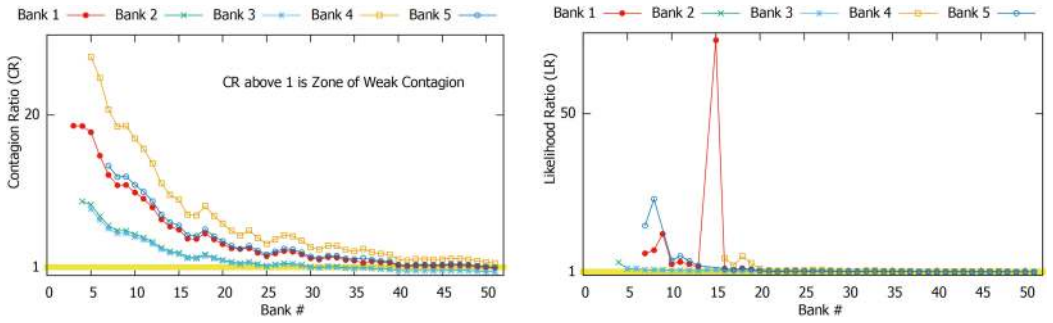

Fig. 2. Performance on real-world data: (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right).

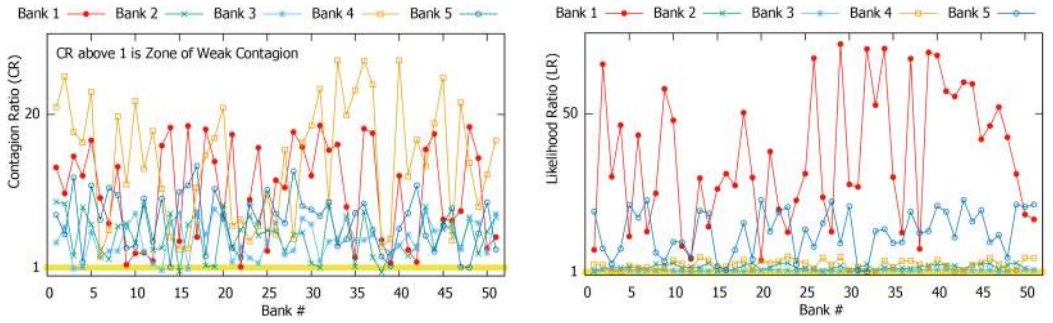

Fig. 3. Performance on Instance #1 of synthetic data: (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right).

We examine the potential for contagion from the failure of the five largest banks (Bank #'s 1-4 in the figures). Taking each of these in turn as the triggering bank, we then take the dysfunctional set D to be consecutive pairs of banks, e.g., the first dysfunctional set under Bank #1 is Bank numbers 2 and 3, the next dysfunctional set consist of Bank numbers 3 and 4, and so on. As peformance metrics, we study 'Contagion Ratio' (CR) and 'Likelihood Ratio' (LR), where we define CR to be the ratio of the LHS of inequality (17) to the RHS of the inequality. We term CR as 'weak' if it is greater than 1. We define LR to be the relative probability of organizational dysfunction through independent direct cyber-shocks and through contagion, calculated as the ratio of the RHS of (18) to LHS. To clearly screen out the pattern of LR and CR values with respect to a varying dysfunctional set, we run additional experiments per triggering node in the set of Bank #s 1 to 4, where for each triggering bank, i, we study the trends for varying β_i and λ_i values for i. We illustrate the results of the study through plots from Figures 9–12.

6.2 Experimental Results

From Figures 2(a)–4(a), we observe that CR is weak for most organizations, validating our theory that cyber-blackouts through strong contagion effects are less likely. CR fails to be weak only when banking organizations in the dysfunction set D are much smaller (in revenue worth) than the triggering bank. Moreover, the value of CR reported for each bank shows how much β_i would have to be to reverse the direction of inequality (17). In this sense, the plots in Figure 2(a) are robust to the estimated values of β_i. Expanding the size of set D makes contagion weaker because of the relative magnitudes of w_i and λ_i^{-1}. High values of LR indicate the dominance of the probability of organizational node dysfunction through independent shocks over node dysfunction through

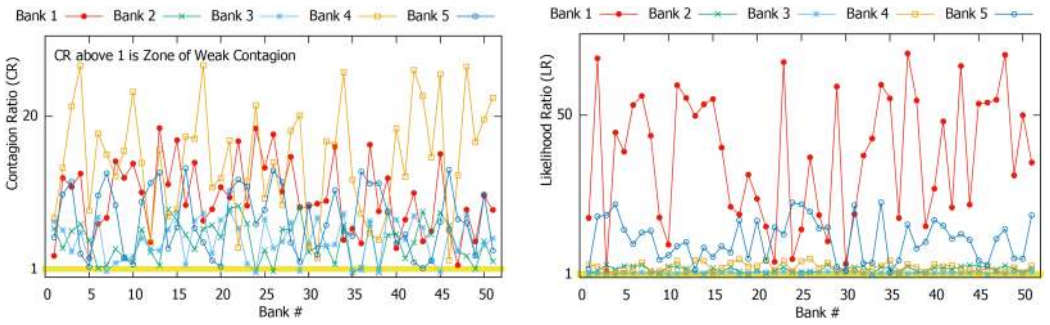

Fig. 4. Performance on Instance #2 of synthetic data: (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right).

contagion. Our plots show that the LR is mostly greater than 1, validating our theory that contagion does not play a major role in organizational dysfunction in the event of a cyber-attack. From Figures 3 and 4 (2 of the 50 random synthetic instances), we observe that network topology does not have a significant role in shaping the contagion phenomenon, i.e., majority of organizations are in the weak contagion zone. However, the CR ratios differ from topology to topology as evident in differences in the plot structure between Figure 2 and Figures 3 and 4. Even if we do not have information on the organizational liability topology of our real-world data, it is evident there is a structure to that topology compared to those characterizing synthetic data. From Figures 9 to 12, we observe that CR values decrease with increasing β_i for each triggering node i —primarily because increasing the fraction of liabilities (β_i) reduces the chances of weak contagion, and high values of β_i leads CR to be less than 1 indicating strong contagion. With increasing size of set D, CR increases reflecting the difficulty to dysfunction a larger set D —a case indicating weak contagion. Given a fixed D not containing triggering node i, with increasing $\lambda_i = \frac{c_i}{w_i}$ (increased leverage) CR decreases indicating reduced chances of weak contagion. With respect to LR, increasing the size of set D increases LR reflecting an increased propensity towards weak contagion—simply because it is difficult to difficult a large set D. Increasing β_i decreases LR, again indicating the negative impact of the higher fraction of liabilities on ensuring weak contagion. The higher the λ_i values, the lesser the value of LR relating to reduced chances of weak contagion—simply because higher λ_i values imply higher cyber-risk for i to become dysfunctional.

Simulations on Non-Bank Settings. Due to lack of real-world data, we are only able to provide simulation results on cyber-blackout scenarios other than the banking sector. Unlike the bank-sector evaluation for which we can estimate graph formation information, we adopt a general (Poisson) random graph formation approach to denote the inter-dependency between organizations. We vary the average degree of a node in the graph from 2 to 8, and keep the liability parameters the same as that of the 2015 dataset, due to the fact that the conceptual essence of liability remains similar across applications. Our rationale to vary the average node degree is due to the the differing nature of applications, that might result in dependency graphs of varying structures. For each average node degree case, we run 50 random instances of graph formation. Due to similarity of the plots (for both, Contagion Ratio and Likelihood Ratio), We showcase a single representative instance of each average node degree case, in Figures 5–8. We observe from Figures 5–8, the similarity of our plots to Figures 1-4. The main message here is the robustness of our results to the network topology, i.e., the network topology does not have a significant role in shaping contagion phenomenon—our inference here is based on simulation data rather than real-world experimental data.

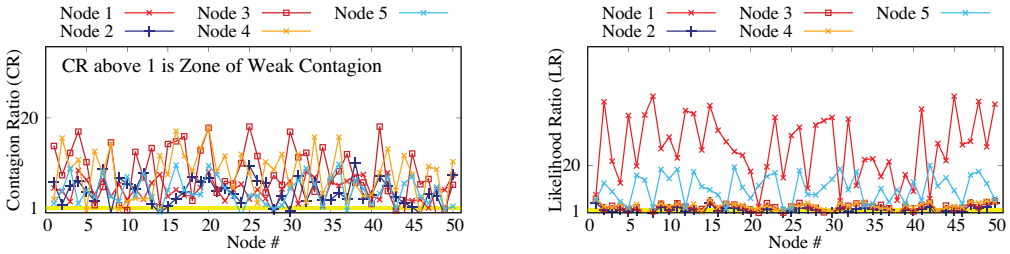

Fig. 5. Performance on instance of synthetic data (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right) **[Case for Average Node Degree 2]**.

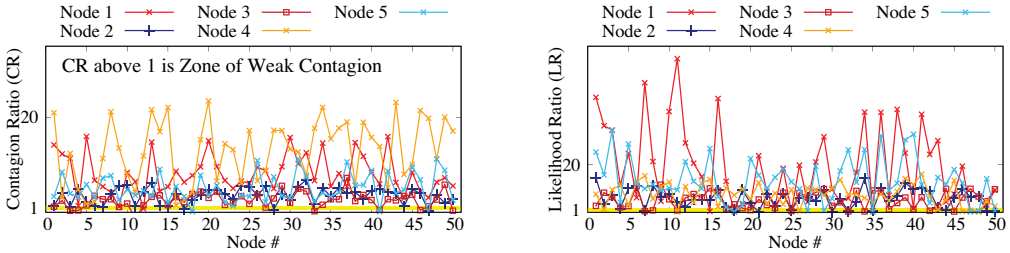

Fig. 6. Performance on instance of synthetic data (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right) **[Case for Average Node Degree 4]**.

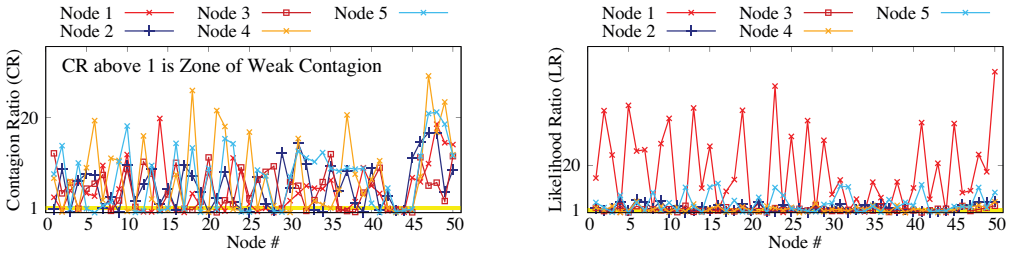

Fig. 7. Performance on instance of synthetic data (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right) **[Case for Average Node Degree 6]**.

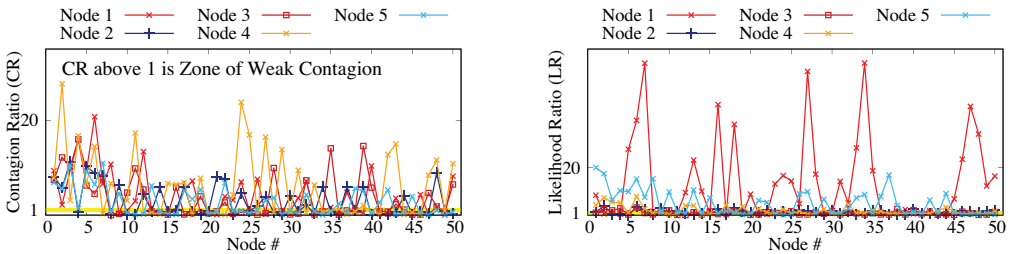

Fig. 8. Performance on instance of synthetic data (a) Contagion Ratio (CR) (left), (b) Likelihood Ratio (LR) (right) **[Case for Average Node Degree 8]**.

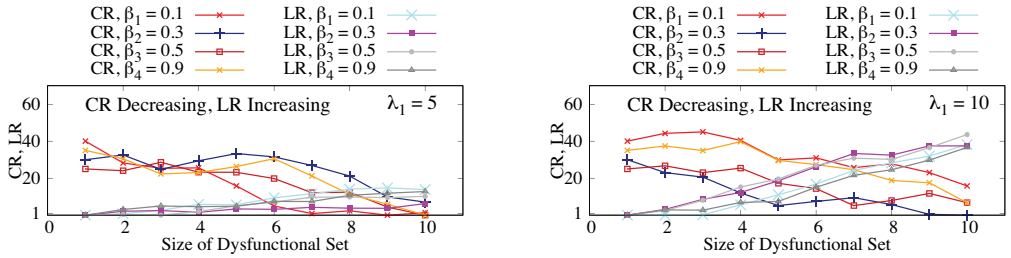

Fig. 9. Performance on instance of synthetic data (Contagion ratio & likelihood ratio) **[Triggering Node #1]**.

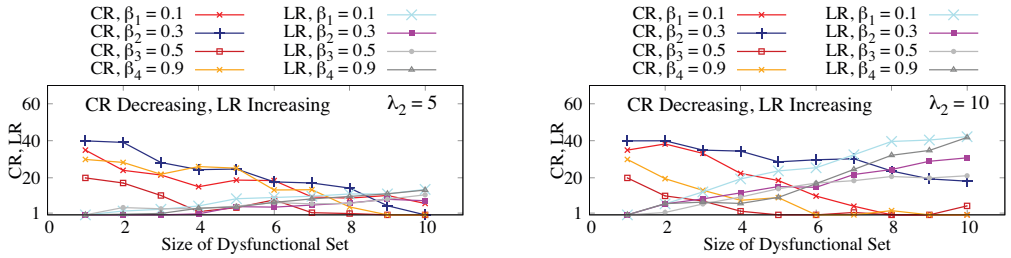

Fig. 10. Performance on instance of synthetic data (contagion ratio & likelihood ratio) **[Triggering Node #2]**.

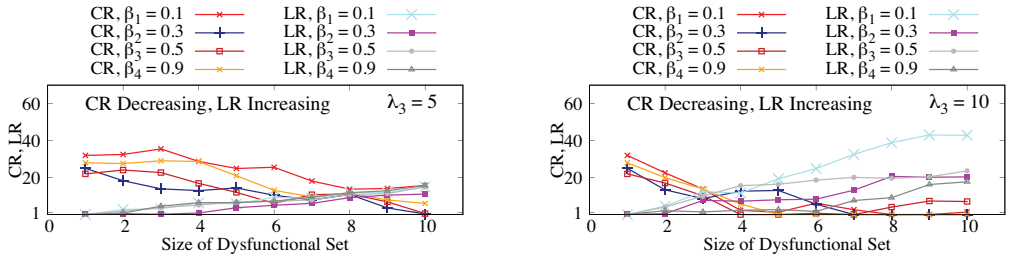

Fig. 11. Performance on instance of synthetic data (contagion ratio & likelihood ratio) **[Triggering Node #3]**.

7 RELATED WORK

In this section, we cite works most related to ours in this article. However, we would like to emphasize upfront that a rigorous analysis (and the consequent impact on the cyber-insurance business) of the likelihood of cyber-blackout phenomena in a network is absent in literature for cyber-insurance or network risk management settings, and our efforts here in this direction are completely new to the best of knowledge. We structure this section in two parts that form a tangential relationship to our work in this article: (a) cyber-insurance market success and (b) risk estimation in network contagion settings.

7.1 Success of Cyber-Insurance Markets

In this work, we investigated worst-case scenarios for a cyber-insurer to cover aggregate cyber-risks. However, a precursor is to have working successful markets in the first place. To this end, recent research works on cyber-insurance [Hoffman 2007; Lelarge and Bolot 2009; Shetty et al.

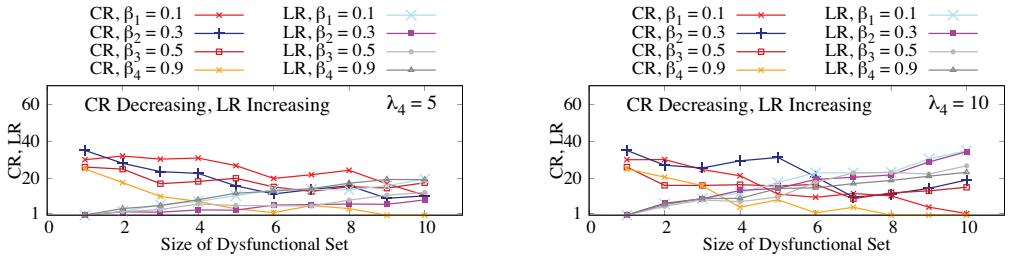

Fig. 12. Performance on instance of synthetic data (contagion ratio & likelihood ratio) **[Triggering Node #4]**.

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 [2009] 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 does it guarantee 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. In addition, 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.

Drawbacks. These works do not investigate issues related to the aggregated risk likelihoods in a networked setting, and their impact on the cyber-insurance industry—a prime determinant for the expansion of the industry.

7.2 Estimation in Attack Spread Settings

In Section 3, we emphasized that evaluating F actually involves mathematically capturing the spread of the infection (attack) vector (e.g., a virus, bot), and is not the focus of this paper. Here, we are only interested in the process of the spread of "organizational dysfunctionality" due to cyber-attacks. The interested reader is referred to [Lelarge and Bolot 2009; Lorenz et al. 2009; Ganesh et al. 2005] to get insights on statistical mean field models to mathematically evaluate F. To the best of our knowledge, no work exists on the spread of "organizational dysfunctionality" due to cyber-attacks as we imply in this article. In terms of the process of the 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], and 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.

- *Drawbacks.* Attack propagation does not imply service disruption. To this end, none of the above works investigate the propagation of service dysfunctionality in a network.

8 DISCUSSION AND SUMMARY

In this section, we first discuss about the current commercial cyber-insurance market and the degree of its inclination towards covering cyber-catastrophes, and follow it up with a summary of this paper. Finally, we peep into future work by discussing an intuition related to current skepticisms behind the statistical feasibility of covering aggregate heavy tailed cyber-risks by a risk management firm. This intuition consequently calls for a deeper analysis into the statistical and economic feasibility of expanding cyber re-insurance markets for catastrophic cyber-risks.

8.1 The Nature of Current Commercial Cyber-Insurance Offerings

Our work in this article has looked into the future of cyber-insurance coverage for the interdependent IT service sector, with respect to quantifying the probability of a cyber-blackout. However, the cyber-blackout scenario though quite relevant for current (and future) general cyber-insurance scenarios (energy, property, marine, aviation, etc.), is not primarily considered, i.e., mostly excluded, while selling insurance policies at present, simply due to profit-minded cyber-insurance agencies being considerably risk-averse on a ruin event arising for correlated and aggregate risk.[3]

To provide support for the above argument (based on data from Coburn et al. [2018]), currently, an estimated half of all cyber insurance policies sold are for limits of less than US$1 million. Limits of over US$10 million are rare (less than 10% of policies written), and for a company to obtain cyber-insurance coverage of US$100 million or more requires the construction of complex 'towers' of coverage involving many different cyber-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. The losses to a company from a (catastrophic) cyber attack can be many hundreds of millions of dollars. The cyber-insurer is providing some financial assistance to its policyholders in the event that they suffer an attack, but is by no means indemnifying their losses as insurers do in other lines of insurance. Companies are left to fund most of the big losses themselves. In general, we estimate that cyber-insurers bear less than 10% of the losses that occur each year. If there were to be a major cyber-catastrophe where large numbers of companies were hit by substantial losses, the cyber-insurance industry would probably bear 15–20% of the total loss experienced by the economy. The cyber-insurers are maintaining their profitability levels, averaging around half of the annual premium generated being paid out in claims, through tightly managed limits and deductibles representing good, safe cyber-risk management. The technique of writing a diversified portfolio of relatively small limits across large numbers of customers is standard practice for spreading the risk.

[3]Warren Buffet and PWC (FTSE Global Markets 2016) have urged against being cautious against the perils of providing cyber-insurance for modern cyber-risks, unless currently privatized and profit-minded markets have serious government intervention to tackle tail-risks in cyber-catastrophe events.

In an emerging market like cyber-risk, where the true nature of the risk is not yet well understood, the cyber-insurers need to 'buy loss experience' building up a database of claims year on year that will help them understand the risk and its characteristics. Subsequently, for cyber-insurance to become a significant-sized market, companies need to be offered limits that are meaningful against the losses that they face. For cyber-insurance companies to offer larger limits, they have to increase the capacity that they make available to cyber-risk. Capacity allocation depends on these insurance companies feeling confident that they have adequately assessed, and priced in, the risk of cyber-catastrophe. A way forward is for cyber-insurers and cyber-re-insurers build their own internal models of cyber-risk, including robust estimates of tail and correlated cyber-risks (such as in Pal et al. [2019]) and costs of risk capital.

Current statistics on cyber attacks show that rate (evaluated using state-of-the-art industry techniques) of cyber-loss vary very significantly for businesses of different sizes, and also between different business sectors. The demand for cyber-insurance, driven by the risk and, more importantly, by the perception of the risk, is similarly varied. A key segmentation is between the insurance market for small and medium-size enterprises (SMEs) and the market for big individual accounts, the large and premier companies. SMEs are a more volume market, with standardized policies and lower premium payments, but tend to have lower cyber-security standards. Big accounts require more customized insurance terms and individual careful (and expensive) underwriting, and are likely to be more targeted by cyber-attackers. The very largest companies (Forbes Global 2000 companies, for example) tend to self-insure, so the big-account insurance market is dominated by large second-tier corporations. Over half of the demand for cyber-insurance comes from companies in the IT, financial services, retail, and healthcare sectors, so it is natural for insurers to end up with concentrations of these in their portfolios. Cyber-re-insurance is a possible option to cover large-valued risks due to a blackout event triggered by a cyber-catastrophe, but for any sorts of reinsurance the risks of the individual policies must be aggregated. In this regard, the regulations affecting the risk of each company would not be treated differently than any other risk that differs across companies / individual policies. In the case of big service-providing companies (e.g., Google), the latter currently (as above-mentioned) do not burden themselves with the risk of those using their services. Cyber-re-insurance services can be complemented via deficit financing methodologies supported by the government agencies. However, steps should be taken to address major geo-political and geo-policy issues that might prevent a wide-spread adoption of deficit financing (an interesting direction for future research work), without which re-insurance services will remain a private venture, built upon some assumptions made in this article.

8.2 Summary of This Article

In this article, we studied the general question: *Is a cyber-blackout in a service organizational network likely?* More specifically, we estimated the probability that all or a major subset of nodes in the network become dysfunctional to provide service in the event of a cyber-attack, a situation which we define as a *cyber-blackout*. The motivation for our research stems from the fact that service liability interconnections among networked IT-driven service organizations create potential channels for cascading service disruptions due to modern cyber-crimes such as DDoS, APT, and ransomware attacks, and cause a bankruptcy-scare effect amongst cyber-insurers via covering aggregate cyber-risk. This scare-effect is the root cause behind cyber-insurers not opening up their coverage capacities enough to boost the cyber-insurance market to prepare for risk management in the age of modern cyber-attacks. As part of our research contributions, we first designed a graph-based model of service obligations, *GSOM*, between organizations in a service chain network. In the event of a cyber-attack, given the values of losses at the nodes in the network, GSOM computes the vector of service valuations that clears the network, and identifies the nodes in the

chain that are dysfunctional to provide service. Using GSOM, we then analyzed (i) how likely it is that a given set of target organizations will become dysfunctional due to contagion from a single source organization, as compared to the likelihood that they become dysfunctional from direct losses to their own service-related assets that does not require dependency on other nodes, and (ii) how much does the underlying network of service dependencies contribute to the increase in the probability of dysfunction of target nodes and corresponding expected value of losses, compared to a situation when there are no network links. As a surprising result, we showed that the loss probability is larger in the absence of network connectivity than that in the presence of network connectivity, implying that simple network spillover effects have a limited impact (except under specific conditions) with respect to service obligations between organizations. We also showed that total additional losses due to network spillover effects are surprisingly small under a wide range of joint distributions for plausible values of model parameters. Finally, we expanded the set of attack sources from a single node to multiple nodes, and studied the negative impact of simultaneous attacks on the entire network. We again showed that the increase in losses due to network interconnections are mostly very small, independent of the network structure and under general assumptions about the joint loss distribution—the primary rationale being attributed to degrees of heterogeneity in wealth base among organizations, and Increasing Failure Rate (IFR) property of loss distributions. Thus, the results obtained through our work encourages cyber-insurers to open up their coverage capacities for a healthy cyber-insurance market.

8.3 Peeping into Future Work

Thus far, we have primarily addressed the effect of the network and mathematically light-tailed IFR cyber-risk distributions on cascading service dysfunctionality phenomena. In reality, modern age catastrophic cyber-loss distributions could be heavy-tailed in nature. In this section, we provide a statistically intuitive explanation behind aggregate cyber-risk managers like re-insurance services being skeptical of resolving the problem of successfully covering aggregate cyber-losses, when an IT-driven liability-networked system is subject to catastrophic cyber-attacks. The crucial importance of this statistical intuition lies in it driving (via its impact on cyber-insurer utility) a formal strategic/economic model, as part of future work, that will help us establish the (in)-effectiveness of re-insurance services provided via a competitive market.

Risk Aggregation due to the Liability between Organizations. Cyber-insurance firms that cover organizational losses generally take on a limited coverage liability. In addition, they insure service organizations that are often liable for maintaining the QoS of firms dependent on the latter. In such situations, though the insurers of individual firms have the advantage of diversifying their coverage to multiple other insurance firms of organizations they depend upon, there is a disadvantage as well that arises when a cyber-attack causes significant cascading losses in a network of service-liable organizations. Intuitively, it is not certain here that a risk-averse cyber-re-insurance firm will always find it profitable (or even feasible) to cover aggregate losses for supply-chain networks of IT-driven industries.

Intuition via Cyber-Risk Distributions. 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,[4] i.e., a subclass of distributions whose 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 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 random variable (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 *decrease* in VaR[5] 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 and cdf are respectively given by

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

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

Let $L_{\mu,\sigma}$ be the class of random variables (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

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

[4]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 the Levy distributions are the only stable distributions for which closed form expressions exist and consequently help in tractable analyses.

[5]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, 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, that 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.

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.

Intuition-Driven Pratical 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 an uniform $\frac{1}{K}$; for each $i \in \{1, \ldots, K\}$, we will have $\sigma = \left(\sum_{i=1}^{K} (c_i \sigma_i)^{\frac{1}{2}}\right)^2$ 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]). 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.

APPENDIX A. GSOM EXAMPLE

In this Appendix, we exemplify the GSOM framework via a graphical example. To illustrate the spread of cyber-attack shock-effects among liability-induced insured networked organizations, consider the example in Figure A.1 (left). The number on each directed edge represents a payment obligation, and each nodes net worth is shown in bold in blue. For example, consider firm C. It is owed 160 currency units (CUs) by outside entities, and it owes 50 CUs to a possibly different set of organizational entities as a result of cyber-attacks in the past. Additionally, C is owed 100 CUs by organization B and it owes 100 CUs to each of organizations A and D. The difference between C's assets (160 + 100) and its liabilities (50 + 100 + 100) result in it's net worth of 10 CUs.

Suppose now that the given organizational network is hit by a cyber-shock that causes some organizations to default on their payments to C: instead of the promised 160 CUs, they pay only 40 CUs. Then C becomes dysfunctional because its assets total 100 + 40 = 140 CUs, whereas it owes 50 CUs to internal losses after the cyber-attack and 200 CUs to other entities, as part of service liability. In this case, we assume that C's remaining assets are paid pro-rata to C's creditors. As we shall see, C's assets may turn out to be worth even less than 140 CUs, because its default may trigger a chain of organizational node dysfunctions that lead back to C.

To work through these spillover effects, we proceed by computing "interim" payoffs as follows: If we take the interim value of C's assets to be 140 CUs, the pro-rata rule implies that C pays (100/250) × 140 = 56 CUs to D, 56 CUs to A, and 28 CUs to itself (reserves for its self non-liable losses). Now D has assets worth 204 + 56 = 260 CUs and debts totalling 300 CUs, so D is in a dysfunctional state. The pro-rata rule implies that D pays 130 CUs to A and 130 CUs to itself. At this stage, A's assets have an interim value of 120 + 130 + 56 = 306 CUs, whereas its nominal obligations come to 360 CUs. Thus, A becomes dysfunctional and the pro-rata rule implies that it pays one-half of its assets to B, namely 153 CUs, and an equal amount to itself. At this juncture, B's assets are worth 153 + 30 = 183 CUs, whereas its obligations total 200 CUs. Therefore, B becomes dysfunctional; it pays 91.5 CUs to C and 91.5 CUs to itself.

At this point, we discover that the value of 140 CUs that we used for C's assets was incorrect. That value reflected the initial outside cyber-shock of 40 CUs, but it assumed full repayment

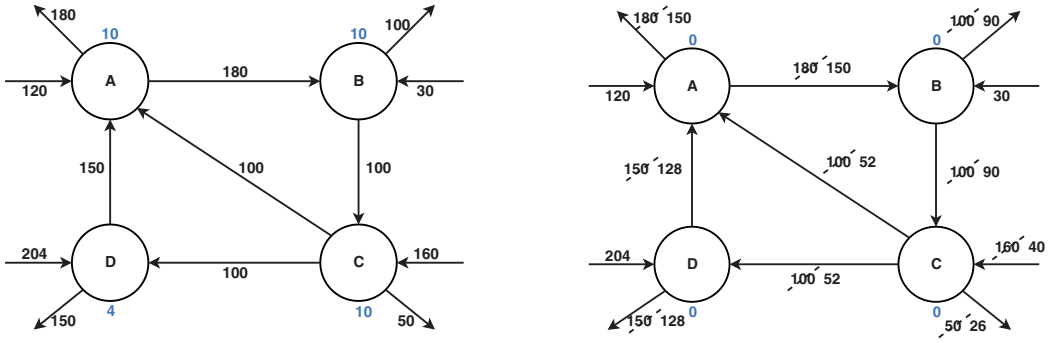

Fig. A.1. GSOM framework: An example of service network with payments due (a) before cyber-shock (left), (b) after cyber-shock (right).

of 100 CUs from B. In fact, B is able to pay at most 91.5 CUs, so C's assets are worth at most 131.5 CUs and the cycle must be repeated. Because of this cascade of node dysfunctions, determining the consequences of the initial cyber-shock is a fixed-point problem. The reader may verify that successive application of the fixed-point algorithm leads to the solution shown in Figure A.1 (right). Here the incoming payments equal the outgoing payments at every node, and the payments from each node are distributed in proportion to the nominal amounts owed. Thus, we obtain a mutually consistent (equilibrium) set of payments; moreover, it is unique.

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.

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 E. Barlow and Frank Proschan. 1975. *Statistical Theory of Reliability and Life Testing: Probability Models*. Technical Report. Florida State Univ Tallahassee.

Richard S. Betterley. 2015. Cyber/privacy insurance market survey-2015. Advisen Annual Report.

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.

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.

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.

Steven N. Durlauf. 1993. Nonergodic economic growth. *The Review of Economic Studies* 60, 2 (1993), 349–366.

Larry Eisenberg and Thomas H. Noe. 2001. Systemic risk in financial systems. *Management Science* 47, 2 (2001), 236–249.

F. J. Fabozzi and H. M. Markowitz. 2002. *The Theory and Practice of Investment Management*. John Wiley & Sons.

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.

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 *24th Annual Joint Conference of the IEEE Computer and Communications Societies. (INFOCOM'05)*, Vol. 2. IEEE, 1455–1466.

Jianxi Gao, Sergey V. Buldyrev, H. Eugene Stanley, and Shlomo Havlin. 2012. Networks formed from interdependent networks. *Nature Physics* 8, 1 (2012), 40.

Paul Glasserman and H. Peyton Young. 2016. Contagion in financial networks. *Journal of Economic Literature* 54, 3 (2016), 779–831.

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. *Wired, June* 20 (2017). https://thehill.com/opinion/technology/380948-when-nation-states-hack-the-private-sector-for-intellectual-property.

Andy Greenberg. 2018. The untold story of NotPetya, the most devastating cyberattack in history. *Wired, August* 22 (2018). https://www.theregister.co.uk/2017/05/12/nhs_hospital_shut_down_due_to_cyber_attack/.

Steve Grobman. 2018. When nation-states hack the private sector for intellectual property. https://thehill.com/opinion/technology/380948-when-nation-states-hack-the-private-sector-for-intellectual-property.

Kat Hall. 2017. UK hospital meltdown after ransomware worm uses NSA vuln to raid IT. https://www.theregister.co.uk/2017/05/12/nhs_hospital_shut_down_due_to_cyber_attack/.

Annette Hoffman. 2007. Internalizing externalities of loss prevention through insurance monopoly. *Geneva Risk and Insurance Review* 32 (2007).

Glyn A. Holton. 2003. *Value-at-Risk*. Academic Press.

Norman L. Johnson, Samuel Kotz, and N. Balakrishnan. 1995. *Continuous Univariate Distributions*, vol. 2 of Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley, New York.

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 La. 2018a. Influence of clustering on cascading failures in interdependent systems. *IEEE Transactions on Network Science and Engineering* (2018).

Richard J. La. 2016. Interdependent security with strategic agents and cascades of infection. *IEEE/ACM Transactions on Networking* 24, 3 (2016), 1378–1391.

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.

Risk Management Solutions, Inc. 2016. *Managing Cyber Insurance Accumulation Risk*.

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

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

Albert W. Marshall and Ingram Olkin. 1974. Majorization in multivariate distributions. *The 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.

Pascal Millaire. 2017. 3 reasons why the insurance industry will never be the same after the Mirai ddos attack. *Advisen's 2017 Cyber Guide*.

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. *The Review of Economic Studies* 67, 1 (2000), 57–78.

Mark Newman. 2018. *Networks*. Oxford University Press.

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.

Ranjan Pal, Leana Golubchik, Konstantinos Psounis, and Tathagatha Bandyopadhyay. 2019. On robust estimated of correlated risks in cyber-insured IT firms: A first look at optimal AI-based estimated under small data. *ACM TMIS* 10, 3 (2019).

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, Konstantinos Psounis, and Pan Hui. 2014. Will cyber-insurance improve network security? A market analysis. In *INFOCOM, 2014 Proceedings IEEE*. 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* (2017).

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.

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

Jordan Robertson and Michael Riley. 2018. The Big Hack: How China Used a Tiny Chip to Infiltrate U.S. Companies. *Bloomberg Businessweek* 4 (2018).

Leonard C. G. Rogers and Luitgard A. M. Veraart. 2013. Failure and rescue in an interbank network. *Management Science* 59, 4 (2013), 882–898.

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.

AWS Sales. 2018. Case Studies & Customer Success - Amazon Web Services (AWS).

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.

Symantec. 2016. Attackers target both large and small businesses. https://seekingalpha.com/article/4224061-why-sold-apple-stock.

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

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.

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

Wikipedia Contributors. 2018. 2007 cyberattacks on Estonia. *Wikipedia* (2018).

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

Received November 2018; revised September 2019; accepted February 2020