Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

CYBER (RE-)INSURANCE POLICY WRITING IS NP-HARD IN IOT SOCIETIES

Ranjan Pal
Taoan Lu
Peihan Liu

Electrical Engineering and Computer Science
University of Michigan
1301 Beal Ave
Ann Arbor, MI 48109, USA

Xinlong Yin

College of Computing
Georgia Institute of Technology
801 Atlantic Dr.
Atlanta, GA 30332, USA

ABSTRACT

The last decade has witnessed steadily growing markets for cyber (re-)insurance products to mitigate residual cyber-risk. In this introductory effort, we prove that underwriting simple cyber re-insurance policies can be worst case computationally hard, i.e., NP-Hard, especially for upcoming IoT societies. More specifically, let alone human underwriters, even a computer cannot compute an *optimal* cyber re-insurance policy in a reasonable amount of time in worst case scenarios. Here, optimality of a contract is judged based on the extent of information asymmetry induced negative externalities it mitigates between a re-insurance seller and a buyer. Our result does not challenge the existence of cyber re-insurance markets that we feel will be a necessity in the IoT age, but only rationalizes why their growth might be slow, and would subsequently need regulatory intervention. *As a direct applicability of our methodology, we argue that optimal traditional cyber-insurance underwriting in IoT societies is also NP-Hard.*

1 INTRODUCTION

IoT-driven smart cities are examples of complex service networked ecosystems that are popularly on the rise around the globe, with major cities like Singapore, Dubai, Barcelona, and Amsterdam being working examples. The proper functioning of such cities, as a system of systems (SoSs), is hugely based on the success of supply chain relationships from diverse, often automated, and critical sectors such as automobiles, electronics, energy, finance, aerospace, etc. In the IoT age, these relationships are often realized via large scale systemic network linkages between organizations, (see Figure 1.1. in (Coburn et al. 2018)), in each of which many salient functionalities rely on the interplay of IoT hardware (e.g., sensors, actuators, cameras), application software (e.g., Oracle for DBMS support, cloud service software), and IoT firmware. With the rapid organizational deployment of the Internet of Things (IoT) in most countries (projected to contribute to a multi-trillion smart-city economy by 2025), and complemented via the use of state-of-the-art data science methodologies, operations in the above-mentioned service sectors are getting increasingly automated and cost-effective. Most importantly, the potent combination of IoT, smartphones, and data science is continuously opening doors to a plethora of novel and pervasive consumer services in each of these sectors that are benefitting the economy and society as a whole.

1.1 Cyber-Security Concerns and The Evergreen Inefficiency of Cyber-security Product Economies

A major flip-side to the significant promise of modern IT/IoT societies in garnering benefits for the socio-economic landscape, is the much-desired but non-existent strong cyber-security that should complement modern IT and IoT infrastructures and the services they provide. Corporate surveys conducted by popular

978-1-6654-3311-2/21/$31.00 ©2021 IEEE

cyber-risk management firms (e.g., Advisen, PartnerRe, Deloitte) keep confirming every year (via their annual reports) that most industries around the globe – small, medium, or big, are successfully breached through malicious events that include cyber-extortion (e.g., ransomware), unintentional data disclosures, lost or stolen data, data breaches, unauthorized data collection and disclosure, identity theft, network/website disruption, business email compromise via social engineering, and denial of service. This, despite the rapid growth in the number of cybersecurity technology solutions over the years. The reason for this strange paradox is multi-fold:

1. Despite the rising trend in the last few years (thanks to multiple major cyber-catastrophes such as Mirai DDoS, WannaCry/Petya ransomware attacks, and Sony/Target data breaches hitting society) among organizational boards acknowledging cyber-risk to be a top five concern for business continuity, reputation, and profitability (Shetty et al. 2018; Pooser et al. 2018; Gatzlaff and McCullough 2010), the proportional time, resources, and effort have not been put in to design effective board-level policies that aptly incentivizes employees to behaviorally improve cybersecurity practices, and make best use of installed security products.
2. Cybersecurity technology solutions is a *market for lemons* (a term coined by economist George Akerlof in 1970 (Akerlof 1978)). The root of the technology efficacy problem is primarily economic driven by information asymmetry between the parties that prevent technology buyers (e.g., the CISOs and the enterprise team of organizations) from effectively evaluating technology, and incentivizing security vendors to sell sub-optimal solutions in the market, that are not as effective as promised and which reduce trust in cybersecurity technology. More specifically, the technology solutions market is too congested with tons of products for buyers to give in quality time and effort to evaluate and rank the effectiveness of each product – to the extent that the buyers believe that lack of quality is the reason behind too many products to exist concurrently in the market.
3. Cybersecurity products are often an outcome of a high-risk “casino economy” where a fragmented vendor industry is configured to manufacture products they think (a) venture capitalists (VCs) will invest in, (b) that larger companies might want to integrate and make the smaller companies follow suite, and (c) that customers can be convinced to buy. These products, akin to a gamble, might be innovative enough to occasionally help cybersecurity, or it might not, but frankly nobody has a clue. Now iterate this process over 10-15 years and you end with a lot of complexity, layers of obsolete and often non-performing technology that requires ever more scarce human expertise to maintain and keep running. As Ciaran Martin, former head of the UK National Security Centre (NCSC) once said “In cybersecurity right now, trust doesn’t always sell, and good security doesn’t always sell and isn’t always easy to buy. That’s a real problem.”
4. The IoT has ushered in a new and difficult challenge to technologically manage cyber-risk in modern IoT societies. Billions of IoT devices (currently, tens of billions, and projected by Cisco to be a whopping 125 billion by 2030) are deployed in most industrial sectors, and connected as part of intra and inter-organization networks. Most of these devices are unattended for long periods of time, have poor security features due to limited computational capabilities – incapable of running sophisticated security tools even if desired, and often loaded with weak default passwords (Gilchrist 2017). This has made Industrial IoT-driven cyber-physical industry systems quite easy to be breached, as evident from the catastrophic Mirai botnet attack a few years ago. The denser the IoT penetration, greater the likelihood of system and systemic cyber-risk, and consequently greater the negative social and economic impact (Coburn et al. 2018; Tanczer et al. 2018).

1.2 The Steady Rise of Cyber-Insurance Markets as a Complement to Cybersecurity Product Markets

Bolstered by our above viewpoints, and already common knowledge since the beginning of the millennium, is the fact that technology products are solely not going to lead us to an ideal cybersecurity state due to them forming an inefficient economy primarily driven by unresolved information asymmetry challenges –

a reality that is not going to change for the better anytime soon. This has led to organizations round the globe embracing cyber-risk management (CRM) solutions that are a mix of both, in-house efforts (e.g., via effectively using security vendor products, raising employee awareness to cyber-security, self-insurance), and commercial third-party cyber-loss coverage products that eliminate residual risk. *One such commercially popular third-party coverage product is cyber-insurance that defends businesses and individuals from the financial losses caused due to cyber risk exposure.* Cyber-insurance covers first-party costs, such as cyber extortion, cyber forensics, credit monitoring, civil fines and penalties, and privacy notification, as well as third party liability such as electronic media liability and network security and privacy liability. *Apart from providing loss coverage as its salient functionality, cyber-insurance carries with it the promise of improving cybersecurity* - known through principle, as well as verified in theory through multiple research efforts since the early 2000s (Anderson and Moore 2009; Lelarge and Bolot 2009; Pal et al. 2014; Pal and Golubchik 2010; N.Shetty et al. 2009; Biener et al. 2015). Today, more IT-driven organizations (ITOs) than ever (nearly 80 percent in the USA) carry cyber insurance, with 55 percent in North America and Europe buying stand-alone cyber-insurance policies (from at least 200 commercial cyber-insurance vendors just in the USA, as reported by Advisen) in an annual market that is worth approximately 8 billion USD globally (projected to grow to 25 billion USD by 2025). Add to this the current push from the legal and policy front in certain parts of the world to invest in cyber-insurance. For example, in February 2020, the Californian assembly introduced a bill to make cyber insurance mandatory to process regulated and protected personal information for all state contractors. The rise in data privacy laws, such as the Personally Identifiable Information (PII) and the Health Insurance Portability and Accountability Act (HIPAA), in the US; the global standard, Payment Card Industry-Data Security Standard (PCI-DSS); and the European Union's (EU) General Data Protection Regulation (GDPR) are persuading insurance providers to focus on cyber insurance measures. In February 2020, the European Insurance and Occupational Pensions Authority (EIOPA) released its strategies for cyber underwriting in order to build a strong cyber insurance market.

1.3 The Need for Cyber Re-Insurance Markets in the IoT Age to Manage Aggregate Cyber-Risk

It is obvious that the ushering pervasive IoT age with 100s of IoT devices per home/organization will bring forth the need for businesses and homes to increasingly buy coverage CRM solutions like cyber-insurance. This is simply because the cyber-attack space will be broad enough in the digital terrain (both, in time and space) for humans to always prevent being security-hacked by smart adversaries. In addition, the high likelihood of correlated cyber-risks (as those that caused Mirai-type cyber-attacks) affecting a fragile network of IoT devices, and super organizational networks formed by them with lead to cascading cyber-risks (due to same device software, hardware, and firmware models being affected at a given time across cyber-space) that will aggregate at CRM solution providers like insurers. *In such scenarios, the idea of effectively spreading aggregate cyber-risk among multiple risk managers like cyber (re)insurers) as a potent aggregate cyber-risk management methodology is gaining traction*[1] (Coburn et al. 2018; Kessler 2014; LI 2018) for IoT-driven smart society settings. Here, insurers covering aggregate cyber-risk of organizations in a given sector (e.g., manufacturing) wish to spread that risk among insurers of firms that are higher up in the supply chain (e.g., energy companies). However (a) it is not a trivial question to formally judge in favor of the effectiveness of this idea, and (b) there may be significant differences in the cyber and non-cyber re-insurance settings, where benefits of non-systemic outcomes in the latter (as qualitatively stated in (Kessler 2014)) may not apply to the former. This is simply due to (a) the unique correlated property of cyber-risks in space and time, and (b) the potentially unavoidable information asymmetry between a re-insurer and its client insurers. *Consequently, without a formal analysis, profit-minded and risk-averse re-insurers (usually traditional insurance organizations) may not have the confidence to scale their service markets at the same rate as per market projections* (Welburn and Strong 2019).

[1] More specifically, in 2019, *SwissRe* estimated the size of the global cyber re-insurance market at USD 4.5 billion, and projected an yearly growth of approximately 20%. (Source - *Cyber re-insurance is the "new normal"*, October, 2020.)

1.4 Research Motivation and Contribution

Research Motivation - In this paper, we are motivated by the inherent computational challenges that information asymmetry between a cyber re-insurer and its clients might pose to the former being able to underwrite optimal contracts, i.e., those ensuring economic efficiency. *The aspect of investigating computational challenges is unique to an IoT societies, simply because of the significantly high scale of entities, i.e., IoT-driven organizations, that will buy cyber-insurance (in light of reasons mentioned in Section 1.1 and 1.2), and consequently the combinatorial number of non-transparent ways in which a particular cyber-insurer might package individual organizational insurance contracts, to offer to a re-insurer*, to buy a re-insurance policy, and strategically hide lemon contracts to their advantage, but disadvantage the re-insurer (see Section 2 for more details).

Research Contribution - *As our main contribution, we show in this work that optimal cyber re-insurance underwriting in smart IoT-driven societies is NP-Hard in the worst case, and that too for the most simple of contracts* **(see Section 3)**. Formally, we show that a simple version of cyber re-insurance underwriting reduces from the *densest k-subgraph problem* in theoretical computer science that is known to be NP-Hard (Khot 2006). In more general terms, optimal underwriting of cyber re-insurance in the presence of information asymmetry (IA) between the buyer and the seller is infeasible even for a computer (unless P = NP), leave alone humans. Here, optimality of a contract is judged based on the extent of IA-induced negative externalities it mitigates between a re-insurance seller and a buyer, i.e., the cyber-insurer. *Our result does not challenge the existence of (security improving) cyber re-insurance markets, that we see growing steadily in the IoT age* **(see Section 4)**. It only emphasizes (a) the inevitable need for policy buyers to embrace non-ideal policy parameters (e.g., unfair premiums), and (b) the role quantum computers could play in future in policy design to be able to compute optimal insurance policy parameters - significantly alleviating the cost of information asymmetry. *As a direct applicability of our proposed hardness reduction, the basic cyber-insurance contract design problem in IoT societies is also NP-Hard.* This is primarily due to the nature of correlated cyber-risk whose sources are spread in a service-networked society in highly non-transparent manner, both to the cyber-insurer, and in some cases to the insured - leading to a two-sided information asymmetry problem that is immensely difficulty to tackle (see Section 4 for an explanation).

1.5 Related Work

We review related work in this section in a concise and brief manner in the interest of space. First and foremost, *ours is the first work of its kind to investigate into the computational tractability aspects of cyber (re)-insurance contract design in either IT and/or IoT driven societies.* The proven potential of cyber-insurance to improve cybersecurity has been mathematically shown in seminal papers (Lelarge and Bolot 2009; N.Shetty et al. 2009; Hoffman 2007; Pal and Golubchik 2010; Pal et al. 2014; Naghizadeh and Liu 2014; Pal et al. 2018), though without reaching market efficiency. However, this has not completely discouraged cyber-insurance providers from increasing their supply of solution products, that is steadily seeing an increase over the years (see Section 4 for a rationale). The current advent of cyber re-insurance solutions is fuelled (since 2017) by the recent massive cyber-attack impacts caused by large-scale DoS (Mirai) and ransomware (WannaCry, Petya) attacks that have led to cascading and aggregate supply-chain organizational claims upon insurers. To this end, *recent theoretical efforts have zoomed in to the statistical nature of loss impact distributions, and their influence on the feasibility (if not optimal profitability) of cyber re-insurance markets.* More specifically, in a series of efforts (Pal et al. 2020; Pal et al. 2020; Pal et al. 2020; Pal et al. 2021), the authors have proved that spreading *catastrophic* heavy-tailed cyber-risks that are identical and independently distributed (i.i.d.), i.e., not tail-dependent, *is not* an effective practice for cyber re-insurers, whereas spreading i.i.d. heavy-tailed cyber-risks that are *not catastrophic* is. While this latter point has long been believed and empirically validated in the cyber-insurance research literature, the former point is a surprising new facet that the authors unravel via theory. In addition, spreading *catastrophic* and *curtailed* heavy-tailed cyber-risks that are (non) identical and independently distributed (i.i.d.), i.e., not tail-dependent, *is not* an effective practice for cyber-reinsurers. *Orthogonal to investigating on the statistical feasibility of cyber re-insurance markets, we investigate on the computational feasibility of underwriting optimal cyber-insurance contracts.*

2 AN INTUITION TO WHY CYBER RE-INSURANCE UNDERWRITING IS HARD

The Setting - Consider a single cyber-insurer in the IoT age interested in buying a cyber re-insurance policy to cover for catastrophic cyber-events. It wants to cover for N contracts, each sold to individual organizational clients, via a certain number of cyber re-insurance policies. Assume *each* **individual contract** (IC) carries with it a *"fairness" quotient* (FQ) that is w.l.o.g. binary, and is either 0 or 1 with probability $\frac{1}{2}$ independently of all others. More specifically, we assign the FQ to be 0 for an individual contract if the cyber-insurer *significantly* (based on pre-set insurer thresholds that is usually different across insurers) *under* or *over* estimates client cyber-risk during audit processes, and 1 otherwise[2]. In other words, an IC has an FQ of 1 if the premiums associated with that contract are not "far enough" from mathematically exact fair premiums that are virtually impossible to derive in practice. Thus, expected fairness for the entire contract bundle is $\frac{N}{2}$. Now suppose that this insurer ex-post over time gathers some "inside" information that an n-sized subset S of the contracts are lemons where the policy holders, due to loose regulations on cyber-information disclosure, are taking advantage of information asymmetry, and are paying for an under-priced contract with probability 1. In this case the value of the expected fairness of the entire contract bundle will be $\frac{N-n}{2}$ with a lemon cost of $\frac{n}{2}$ incurred by the insurer.

Why Insurance Derivatives 'Seem' to Reduce Lemon Cost - In principle, the cyber-insurance company, in order to hedge their risk for large N in IoT age, can buy insurance derivatives with a cyber re-insurer to significantly ameliorate the lemon cost (*source* - NAIC). In particular consider the setting where the insurer wants to buy M new **re-insurance claim contracts** (RCCs) from the re-insurer (akin to an insurance derivative in the finance industry), each of them depending on the performance of D bundled (akin to a portfolio) underlying individual contracts (ICs), where some of these contracts may overlap in the portfolio of multiple RCCs *(synonymous with claim contracts henceforth).* A good performance indicates the re-insurer covering less aggregate risk arising from D at any given time instant. *It is usual in practice to have $M \cdot D$ much greater than N to ensure that each IC is packaged in multiple insurance derivative claim contracts.* Assume each of the M contracts generate an expected fairness quotient (EFQ) $\frac{N}{3M}$ to the cyber-insurer as long as the number of ICs that are lemons in an RCC is at most $\frac{D}{2} + t\sqrt{D}$ for some parameter $t = O(\sqrt{\log D})$, and otherwise results in an EFQ of zero. Thus, in the best case, if there are no lemon ICs then the combined EFQ of these M claim contracts is very close to $\frac{N}{3}$ (the ideal case should yield an EFQ of N). One can rationalize using the central limit theorem (CLT) that if the pooling is done randomly (each contract depends on D random ICs), then even if there are n lemon ICs, the value is still $\frac{N}{3} - o(n)$, no matter where these lemon ICs are. More specifically, according to the CLT, the total number of lemon ICs may be assumed to be distributed like a Gaussian. Thus, so long as the fraction of lemon classes is much smaller than the safety margin of t standard deviations, the probability for a single claim contract containing many ICs to generate a significantly low EFQ, is tiny. Thus, we clearly see that insurance derivatives do indeed help significantly to reduce the lemon cost from $O(n)$ to $o(n)$.

Power of Derivatives Dampened by Computational Challenges of the Re-Insurer - The cyber-insurance firm has no incentive to do the above-mentioned pooling completely randomly because he knows S, the set of lemon ICs. Some simple calculations suggest that the optimal strategy for a cyber-insurer is to pick some m of the RCCs, and make sure that the lemon ICs are over-represented in them — to an extent about the scale of $\sqrt{D}$ that is just enough to skew the probability that the claim contracts will in total not result in a zero EFQ for the re-insurer. This is rationalized by the fact that since ICs contained in the same RCC come from different organizations, the yields of non-lemon ICs are uniformly i.i.d., so the expected number

[2] An individual contract may or may not have deductibles associated with it, but we consider EFQs oblivious of the contract type.

of zero-valued FQs among D non-lemon ICs is $\frac{D}{2}$ with variance D. In the ideal case, *a fully rational* (i.e., computationally unbounded) cyber re-insurer, i.e., the 'buyer' (seller) of claim requests (contracts) from the cyber-insurer, can enumerate over all possible n-sized subsets of $[N]$ to verify that none of the lemon ICs are over-represented, thereby upper bounding a lemon cost of $o(n)$. *However, in real-life, the cyber re-insurer is computationally bounded, and hence this enumeration is infeasible for large-sized* $[N]$, as in IoT societies. Put in another way, the cyber-insurer can "plant" a set S of over-represented lemon ICs in claim contracts in a way such that the resulting pooling will be computationally indistinguishable from a random pooling. Consequently, the lemon cost for polynomial time cyber re-insurers can be much larger than $O(n)$ in the worst case - *thereby nullifying the vision that introducing insurance derivatives in cyber-settings will mitigate the lemon cost. In practice, contrary to logic, they will instead amplify it.*

Even Tranching Does Not Alleviate Computational Challenges - In seminal papers, DeMarzo (DeMarzo 2005), and DeMarzo and Duffie (DeMarzo and Duffie 1999) introduced the concept of tranching that takes advantage of signalling in economic theory, and for our setting, promises to mitigate information asymmetry between a cyber re-insurance buyer and its seller. More specifically, DeMarzo and DeMarzo et.al., show (when adapted to our setting) that it is optimal for the cyber-insurer to first bundle ICs and then tranche them in a *single claim contract.* The cyber-insurer can offer the re-insurer the less riskier senior tranche to provide coverage for, and can retain the comparatively more riskier junior tranch. The proportion sold versus retained acts as a signalling[3] mechanism to the re-insurer on the quality of ICs to be re-insured; leads to better re-insurance pricing of ICs; and the lemon costs significantly diminish for the re-insurer. However, there is a catch - *in our IoT society setting, the cyber-insurer is offering M (likely large) RCCs instead of a single one.* DeMarzo et.al.'s analysis has no obvious extension to this case because the potentially renewed signalling mechanism is far more complex than in the case of a single claim contract, where all ICs have to be bundled into a single pool. For a perfectly rational re-insurer capable of exponential time computations, lemon costs do get ameliorated by M RCCs. Precisely, the cyber-insurer randomly distributes ICs into M equal sized pools and defines the senior tranche identically in all of them. In doing so, its signal to the re-insurance policy seller consists of the partition of ICs into pools, and the threshold that defines the senior tranche. *Although, for a perfectly rational re-insurer, this signal turns out to contain enough information, it is computationally intractable for a practical boundedly rational cyber re-insurer to decipher.*

3 FORMALLY VERIFYING THE INTUITION

In this section, we formally verify our intuition that deploying optimal cyber re-insurance claim contracts (those mitigating IA between the insurer and the insured) in large scale cyber-settings such as IoT societies will be quite costly for a cyber re-insurer. In other words, we will prove that the lemon costs for a cyber re-insurer to underwrite optimal re-insurance contracts in IoT societies will be amplified, even in the presence of derivative-centric signaling mechanisms proposed by DeMarzo et.al., that have been practically successful in reducing IA in non traditional non-cyber insurance settings.

3.1 Formal Building Blocks and Proving Cyber Re-Insurance Underwriting is NP-Hard

We describe the formal building blocks that (a) first lead us in this section to formally proving that optimal cyber re-insurance underwriting is NP-Hard, and (b) prepare us to derive tight lemon costs for a boundedly rational cyber re-insurer in the next part of the section.

Formulating Lemon Cost - Consider a cyber-insurer claim contract (akin an insurance derivative) with performance metric FQ defined on N individual contract (IC) inputs. Given the input distribution X over $\{0,1\}^N$, and $n \leq N$, we define the lemon cost of FQ comprising n lemon ICs as

$$\Delta(n) = \Delta_{FQ,X}(n) = \mathbb{E}[FQ(X)] - \min_{S \subseteq [N], |S|=n} \mathbb{E}[FQ(X)|X_i = 0; \forall i \in S],$$

[3] The cyber-insurer knows the identity of lemons, but re-insurance sellers only know the prior distribution of lemon ICs.

where the min operator takes into account all possible ways in which the cyber-insurer could "position" the lemon ICs among the N total ICs while designing its RCC. This lemon cost captures the inefficiency introduced in the market due to the existence of lemon ICs.

The Densest Subgraph Problem - Consider an (M,N,D) bipartite graph with M vertices on one "top" side representing claim contracts by the cyber-insurer, and N vertices on the "bottom" side representing the ICs owned by the cyber-insurer. Let each RCC have an outdegree of D indicating the number of ICs that are packaged with a claim contract. We say that such an (M,N,D) graph G contains an (m,n,d) graph H, if one can identify m top vertices, i.e., claim contracts, and n bottom vertices, i.e., ICs, of G with the vertices of H in a way that all of the edges of H will be present in G. Now fix the parameters in the tuple (M,N,D,m,n,d). The *densest subgraph decision problem* (Khot 2006) for these parameters is to distinguish between the two distributions $\mathscr{R}$ and $\mathscr{D}$ on (M,N,D) graphs, where (a) $\mathscr{R}$ results from choosing for every claim contract (top vertex), D random individual contract neighbors on the bottom, (b) $\mathscr{P}$ results by first choosing $S \subset [N]$ and $T \subseteq [M]$ in a manner such that $|S| = n$, $|T| = m$, and then choosing D random IC neighbors for every vertex outside of T, and $D-d$ random neighbors for every vertex in T, and (c) a choice of an additional d random neighbors in S for every vertex in T.

The NP-Hard Nature of the Densest Subgraph Problem - The computational intractability, i.e., NP-Hardness, of the densest k-subgraph problem was proved (via reduction from k-CLIQUE (T.H.Cormen et al.)) in (Applebaum et al. 2010; Bhaskara et al. 2010), where it is formally stated that given (M,N,D,m,n,d) such that $N = o(MD)$ and $(\frac{md^2}{n})^2 = o(\frac{MD^2}{N})$, there does not exist any $\varepsilon > 0$ and a poly-time algorithm that distinguishes between $\mathscr{R}$ and $\mathscr{P}$ with advantage ε.

Re-Insurer Lemon Cost when Cyber-Insurer Has No Inside Information - Consider N ICs that are distributed i.i.d. each of which has a probability 1/2 bearing an FQ of zero and probability 1/2 of bearing an FQ of 1. The fact that each IC is either a lemon or non-lemon with equal probability is not considered inside information. It is usual in practice for cyber-insurers to assume that some ICs will be lemons in view of traditional IA issues between itself and its clients. In our setting the cyber-insurer requests M claim contracts (akin to insurance derivatives) from a cyber re-insurer, where the expected FQ of each claim contract is based on the D ICs forming the portfolio of the contract. Assume a threshold value $b < \frac{B}{2}$, such that each claim contract has EFQ of 0 if more than $\frac{D+b}{2}$ of the ICs contained in it are lemons, and an EFQ value of $V = \frac{D-b}{2D}\frac{N}{M}$, otherwise. This is similar to a binary valuation setting, and represents claim contracts of the simplest type. Given each claim contract depends on D independent ICs, the number of lemon ICs for each RCC closely follows a Gaussian distribution as D gets larger.

The Densest Subgraph Relating to an 'Adversarial' Insurer Having Inside Information - In the case when the cyber-insurer has inside information with probability 1 that there are n lemon ICs, an adversarial cyber-insurer can carefully design the (M,N,D) graph to increase its return from an RCC policy. Note that though each RCC packages D ICs for its re-insurance portfolio, in order to substantially increase its return on the claim, it suffices for the 'adversarial' cyber-insurer to fix about $\sigma \simeq \sqrt{D}$ of the underlying ICs. More precisely, if t of the ICs contained in an RCC are lemons, then the expected number of lemon ICs in the RCC is $D+t$, while the standard deviation becomes $\frac{\sqrt{D}}{2}$. Thus, the probability of this claim contract resulting in an EFQ of 0 is $\frac{\Phi(t-b)}{2\sigma}$ which starts getting larger as t increases. This further implies that means that the difference in return between an IC pool of zero lemons versus that of t lemons is about $V \cdot \frac{\Phi(t-b)}{2\sigma}$. Now say the cyber-insurer allocates t_i of the lemon ICs to the ith RCC. Given that each of n lemon ICs are contained in $\frac{MD}{N}$ claim contracts, we have $\sum_{i=1}^{M} t_i = \frac{nMD}{N}$, that results in a lemon cost of $V \cdot \sum_{i=1}^{M} \frac{\Phi(t-b)}{2\sigma}$ to the cyber re-insurer. The function $\frac{\Phi(t-b)}{2\sigma}$ being concave for $t < b$, and convex otherwise, the optimal strategy for the 'adversarial' cyber-insurer will have t_i's to be either 0 or $k'\sqrt{D}$, for a small constant k'. Put in other words, the lemon cost for the cyber re-insurer is maximized by choosing some m RCCs, letting each of them have at least $d = k'\sqrt{D}$ edges from the set of lemon ICs. *In (M,N,D) bipartite graph, this property corresponds to an* ***(m, n, d) dense subgraph*** *- representing a set of insurer-manipulated RCCs and a set of lemon ICs that have more edges between them than expected.* However, given that the densest subgraph problem is NP-Hard, unless P = NP, it is not computationally tractable for a computer, forget a boundedly rational cyber re-insurer, to decide whether an an (m,n,d) dense subgraph is embedded in an (M,N,D) graph. **Hence, the optimal cyber re-insurance underwriting problem becomes NP-Hard.**

3.2 Results on Formalizing Lemon Cost for a Boundedly Rational Cyber Re-Insurer

Thus far, we have established the NP-Hardness of the optimal cyber re-insurance underwriting process. In this section, we first show that there indeed exists a set of parameters (m,n,d) for which there is a very small likelihood of a dense subgraph existing in an (M,N,D) graph. Though a cyber re-insurer can verify this case, but being boundedly rational, the main question is at what cost? That brings us to the second result on deriving the costs to do such a verification. We now have our first result stating the condition for the *non-existence of a dense subgraph* that a cyber insurer can embed in IC pools - a necessary condition to prevent lemon costs getting amplified for a cyber re-insurer.

Theorem 1 *Consider a cyber-insurer requesting M re-insurance claim contract policies (acting as insurance derivatives) from a cyber re-insurer, each of them having with D individual contracts (ICs) in its portfolio, where these D ICs are are sampled over an universe of N individual contracts (ICs), and are potentially overlapping on multiple claim contracts. With a high probability, there exists no dense subgraph* (m,n,d) *in a random* (M,N,D) *graph when* $n \ll md, \frac{dN}{Dn} > (N+M)^{\varepsilon}$ *for some constant* ε.

Proof. Let $X_{i,j}$ be the random variable indicating the existence of an edge between IC i and claim contract (insurance derivative) j. For simplicity, we assume the $X_{i,j}$s are independent random variables that with probability $\frac{D}{N}$, take up a value of 1. In the real-world, the $X_{i,j}$s are negatively associated statistically - as a result, our independence-induced results in this paper will more likely hold true in realistic settings. Based on proof methods introduced in (Arora et al. 2011), pick any set A of n ICs, and a set B of m claim contracts (a proxy for insurance derivatives) of size m. The probability there exists atleast md edges between A and B is the probability of $X = \sum_{i \in A, j \in B} X_{i,j} \geq md$. Given the independence assumption on the $X_{i,j}$s, and the fact that sum X has expectation $\mu = \frac{mnD}{N}$, using the Chernoff bound, we have $\Pr[X > (1+\delta)\mu] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu}$, where $\mu = \frac{mnD}{N}, 1+\delta = \frac{md}{\mu} = \frac{dN}{Dn}$, the probability is at most $\Pr[X > (1+\delta)\mu] \leq \left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu} \leq (N+M)^{-\varepsilon md}$. The number of such sets is $\left(\begin{array}{c} N \\ n \end{array}\right)\left(\begin{array}{c} M \\ m \end{array}\right) \leq (N+M)^{n+m}$. As a result, via the use of the union bound, the probability that there exists a pair of sets that has atleast md edges between them is at most $(N+M)^{n+m-md}$, which is much smaller than 1 by the assumption that $n \ll md$. Therefore, a random graph will not have dense subgraphs embedded in it, in high likelihood. ■

Practical Implication - The theorem states that there exists a specific set of parameters m,n,d and an associated relation between them for which the ‘adversarial’ cyber-insurer cannot embed any dense subgraph in an IC pool. In other words, *there is a possibility for the cyber re-insurer to verify in the affirmative that lemon ICs are not over-represented in any IC pool, but the question is at what cost? (see Theorem 2).* On an orthogonal note, a thing worth mentioning is that in the real-world, the cyber-insurer tries to put together a diversified portfolio of ICs that are sufficiently independent, and we capture this fact using a random graph - though, the cyber-insurers, in practice, need not construct a random graph in the industry. Having established the possibility of the non-existence of a dense subgraph in a random (M,N,D) graph, the big question then becomes: *how much lemon cost will a boundedly rational cyber re-insurer accrue to verify the non-existence of a dense subgraph?* We now have the following theorem in this regard.

Theorem 2 *Consider a cyber-insurer requesting M RCC policies from a cyber re-insurer, each of them having with D individual contracts (ICs) in its portfolio, where these D ICs are are sampled over an universe of N individual contracts (ICs), and are potentially overlapping on multiple claim contracts. When* $d-b > 3\sqrt{D}$, *for a given derivative threshold b, and* $n/N \ll d/D$, *an* (m,n,d) *subgraph will generate an extra lemon cost that is at least* $(1-2p-o(1))mV \approx n\sqrt{N/M}$.

Proof. Given a cyber-insurer's claim contract (insurance derivative) manipulation activity, for each such contract it manipulates, let Y be the number of lemon ICs. We know that there are d ICs that come from the set of lemon ICs, each of these d ICs will always have an FQ of 0, and the expectation of Y is $\mathrm{E}[Y] = \frac{D+d}{2}$. The Gaussian approximation holds for large enough D, leading to $\Pr\left[Y \geq \frac{D+b}{2}\right] = 1-p$. In line with proof insights introduced in (Arora et al. 2011), for claim contracts that the cyber-insurer does not manipulate with, we assume the expected number of lemon ICs is x, which then leads x to satisfy $m \cdot \frac{D+d}{2} + (M-m) \cdot x = \frac{N+n}{2N} \cdot \frac{MD}{N}$. This relation holds true as both, the LHS, and the RHS are quantities reflecting the expected number of lemon ICs. Given that $x \geq \frac{D}{2} + \frac{n}{2N} \ \frac{md}{2M-2m}$, the probability that the number of lemon ICs is more than $\frac{D+b}{2}$ is at least $\Phi(-3- \ \frac{md}{2(M-m)D}\Big) = p - \Phi'(-3)\frac{md}{2(M-m)D} = p - O\left(\frac{md}{2(M-m)D}\right)$. The expected count of non-manipulated claim contracts that gives no return is at least $p(M-m) - O\left(\frac{md}{2D}\right) = p(M-m) - o(m) = pM + (1-2p)m - o(m)$. This quantity is $(1-2p-o(1))$ smaller than the expectation without dense subgraphs. Hence, the extra lemon cost incurred by the cyber re-insurer is $(1-2p-o(1))mV$. ■

Practical Implication - In the worst adversarial case when (a) $M \ll N \ll M\sqrt{D}$, (b) $m = \Theta(n\sqrt{\frac{M}{N}})$, in which case a random (M,N,D) graph with an embedded (m,n,d) subgraph remains indistinguishable from a random graph under the densest subgraph assumption, and (c) $b = 2\sigma\sqrt{\log\frac{MD}{N}}$, the theorem implies that a boundedly rational cyber re-insurer will incur a lemon cost of $n\frac{N}{M} = \omega(n)$ to verify the non-existence of a dense subgraph in a random (M,N,D) graph, whereas a perfectly rational cyber-insurer will only incur a lemon cost of $n\frac{N}{2M\sqrt{D}} = o(n)$. It is evident that the lemon cost for a boundedly rational cyber-insurer is orders of magnitude higher, when compared to a perfectly rational one.

4 DISCUSSION

In this section, we (i) argue a direct application of the reduction of the densest subgraph problem to optimal underwriting of traditional cyber-insurance contracts - thereby inferring that plain cyber-insurance underwriting is NP-Hard, and (b) briefly review the incentives for various stakeholders in the cyber (re)-insurance market to reasonably contribute to the increasing density of such markets in the IoT age, despite negative results on the existence of efficiency in such markets.

4.1 Even Cyber-Insurance Underwriting is NP-Hard

The crux of the densest subgraph problem reduction being applicable to the cyber-insurance setting lies in the double information asymmetry problem between the cyber-insurers and their clients, i.e., organizations and/or individuals. Note that the cyber-posture of any organization (and individuals) is dependent on both, visible parameters (e.g., strength of security settings, humans in the loop) as well as latent factors (e.g., knowledge of zero-day attack vector, un-patched OS vulnerability by vendor) unknown to the organization and insurers. Both, visible parameters and latent factors lead to correlated cyber-attacks on multiple organizations service-networked amongst each other. It is clear in view of two primary reasons that a cyber-insurer will not be able, in polynomial time, to enumerate all possible setting combinations of security parameters that might render a particular client or a subset of clients as *lemons*: (i) enumerating all security loopholes in any computer system, be it a single machine, or a group of machines in a networked or distributed setting, is a well known *Turing Undecidable* problem (even harder than an NP-Hard problem) (Pfleeger and Cunningham 2010), (ii) in the absence of strong regulations regarding mandatory cyber information disclosures by organizations, combined with the existence of numerous as-of-yet undetected vendor software vulnerabilities unknown to both the insurer and the insureds, it is computationally infeasible to confidently identify lemon clients by any cyber-insurance agency. *(i) and (ii) together directly transform the optimal cyber-insurance underwriting problem into the same bipartite graph structure as the optimal cyber re-insurance underwriting problem, and consequently the densest k-subgraph reduction applies* - hence rendering the optimal cyber-insurance contract underwriting problem NP-Hard, similar to its re-insurance counterpart. This justifies why there is a wide gap (approximately 350 billion USD annually) between supply and demand for commercial cyber-insurance products, i.e., the risk-averse insurers are not confident enough to design profitable contracts due to information asymmetry challenges.

4.2 Being Reasonably Practical in the IoT Age is Better Than Being Perfectly Rational

We argue that it is reasonable from a viewpoint w.r.t. each major stakeholder in a cyber-insurance market ecosystem to increasingly invest in cyber-insurance, despite the utopic nature of writing insurance and re-insurance contracts. In this work, we identify organizations, cyber (re)-insurers, security product vendors, and regulators (e.g., the government) as the primary stakeholders today in the cyber insurance market.

Organizations - It has been a decade-old and constant C-suite level concern in multiple SMB organizational boards to mitigate reputational damage due to cyber-breaches. Reputational risk can pose a danger to the survival of the largest and best-run businesses by wiping out millions or billions of dollars in market capitalization or future profits. According to IBM, compared to other cyber-risk management costs, by 2020 lost business became the largest contributing cost factor post a data breach for SMB organizational, accounting for nearly 40% of the average total cost of a data breach. Given, (i) inevitable cyber-information asymmetry between the insurer and the insured, (ii) large scale IoT networks under operation posing significant cyber-risk correlations in time and space, and (iii) issues related to humans (organizational employees) in the loop, *it is always reasonable for organizations to invest in cyber-insurance when doing a trade-off between buying a non-ideal policy and risking business customers.*

Cyber-Insurers - The IoT age will test the risk-aversion of most cyber-insurers. Network externalities induced by the omnipresence of software vulnerabilities in a few operating systems (OSs), application programs, and security products, but those that are commonly used by most IT systems around the world, result in correlated cyber-risk threats of significant amounts to the coverage dislike of risk-averse insurers. The likelihood of such statistically non-independent risks increase multi-fold in the current IoT age where billions of devices with poor cybersecurity postures (e.g., default passwords, un-encrypted firmware access, unauthorized backdoor access, lack of use of Secure Socket Layer (SSL) technology to connect IoT to the cloud) are connected with one another to the drooling delight of cyber-hackers ever-ready to launch simple attacks that result in cascading catastrophes (such as in the case of Mirai and WannaCry attacks), leave alone the need for sophisticated ones. Such environments will contribute to major pain points for profit-minded cyber-insurers to be risk-averse enough to underwriting and selling enough policies for the social good. *However, given the high risk-aversion of organizations to protect their reputation, cyber-insurance firms will garner a huge opportunity cost upon themselves not to be able to densify their product markets by taking advantage of the former's high degree of risk aversion, even if through non-ideal contract design.*

Security Product Vendors - A previously discussed fact that a cybersecurity product market is one of lemons, should reasonably favor an increasingly dense cyber-insurance market. *The extremely congested market of cybersecurity solution products is indicative of the fact that no few solutions stand out in price-quality trade off, and there is no incentive for an incumbent product to stamp authority over existing solutions.* Assuming this common knowledge among the CISOs of SMB organizations, diligently investing in cyber-insurance is a wise corporate cyber-risk management strategy.

Regulators and Cybersecurity Interests - *Aware of the socio-geo-political barriers to enable the deployment of strong cyber-information disclosure laws, and thereby keep organizational cyber-posture behavior under check, regulatory bodies such as national and state governments should ideally find it reasonable to support commercial cyber-insurance markets.* This is simply because of (a) the cyber-security improving potential of insurance products is of social importance and interest to regulators, and (b) insurers could, via agreements, share cyber-posture data privately with regulators for better design of targeted cyber-security intervention policies. Overall, the regulators can enforce policies that densify cyber-insurance and cyber re-insurance markets because they would contribute to organizations voluntarily sharing appropriate security liability among themselves to the best extent possible, thereby improving global cybersecurity.

5 STATUS QUO IN CYBER (RE-)INSURANCE MARKETS, SUMMARY, AND FUTURE WORK

The IoT age could not be better for new and existing cyber (re-)insurers. On the whole cyber (re-)insurance has performed well to date (cyber as a class is extremely reliant on reinsurance capacity, with year 2020 seeing circa 40% of cyber premium being ceded to re-insurers) and confidence in the class is growing, thanks to corporations' increasing willingness to invest in cyber (re-)insurance. Such is the opportunity that cyber-insurance market leaders such as *Munich Re* has stated it will seek to double its cyber portfolio by the end of 2025. While this sounds astronomical, even if every re-insurer followed suit, market demand would likely still outweigh supply. This is primarily due to fears related to the magnitude of cyber-risk aggregation that have led senior stakeholders of cyber-insurance firms to have an appetite cautious enough, only to deploy modest capital in this business sector that has subsequently led to its limited growth rate. *As summary*, in this paper, orthogonal to existing efforts that have investigated the feasibility of cyber re-insurance markets for IoT societies on the statistical dimension, we investigated the feasibility of cyber (re)-insurance markets from the dimension of computational tractability of designing optimal cyber re-insurance contracts under information asymmetry. We proved that underwriting simple cyber (re)-insurance policies can be worst case computationally hard, i.e., NP-Hard, in IoT societies, through reduction from the densest subgraph problem in theoretical computer science. Our result does not challenge the existence of cyber (re)-insurance markets, that we see, and quite reasonably, growing slowly and steadily in the IoT age. Rather, it only rationalizes why their growth might be slow, changing which would need regulatory intervention. It also emphasizes the inevitable need for policy buyers today to voluntarily embrace non-ideal policies (catalyzed by information asymmetry challenges) to mitigate risk as a priority. *As future work, we wish to design polytime computable cyber (re-)insurance contracts.*

ACKNOWLEDGMENTS

The authors are supported by the NSF under grants CNS-1616575, CNS-1939006, and ARO W911NF1810208.

REFERENCES

Akerlof, G. A. 1978. "The market for "lemons": Quality uncertainty and the market mechanism". In *Uncertainty in economics*, 235–251. Academic Press.

Anderson, R., and T. Moore. 2009. "Information Security: Where Computer Science, Economics and Psychology Meet". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367(1898):2717–2727.

Applebaum, B., B. Barak, and A. Wigderson. 2010. "Public-key cryptography from different assumptions". In *Proceedings of the forty-second ACM symposium on Theory of computing*. June 5th-8th, Cambridge, U.S., 171–180.

Arora, S., B. Barak, M. Brunnermeier, and R. Ge. 2011. "Computational complexity and information asymmetry in financial products". *Communications of the ACM* 54(5):101–107.

Bhaskara, A., M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan. 2010. "Detecting high log-densities: an O (n 1/4) approximation for densest k-subgraph". In *Proceedings of the forty-second ACM symposium on Theory of computing*. June 5th-8th, Cambridge, U.S., 201–210.

Biener, C., M. Eling, and J. H. Wirfs. 2015. "Insurability of Cyber Risk: An Empirical Analysis". *The Geneva Papers on Risk and Insurance-Issues and Practice* 40(1):131–158.

Coburn, A., E. Leverett, and G. Woo. 2018. *Solving Cyber Risk: Protecting Your Company and Society*. Hoboken, New Jersey: Wiley.

DeMarzo, P., and D. Duffie. 1999. "A liquidity-based model of security design". *Econometrica* 67(1):65–99.

DeMarzo, P. M. 2005. "The pooling and tranching of securities: A model of informed intermediation". *The Review of Financial Studies* 18(1):1–35.

Gatzlaff, K. M., and K. A. McCullough. 2010. "The Effect of Data Breaches on Shareholder Wealth". *Risk Management and Insurance Review* 13(1):61–83.

Gilchrist, A. 2017. *IoT Security Issues*. Boston, Massachusetts: Walter de Gruyter GmbH & Co KG.

Hoffman, A. 2007. "Internalizing Externalities of Loss Prevention Through Insurance Monopoly". *Geneva Risk and Insurance Review* 32.

Kessler, D. 2014. "Why (Re) Insurance is not Systemic". *Journal of Risk and Insurance* 81(3):477–488.

Khot, S. 2006. "Ruling out PTAS for graph min-bisection, dense k-subgraph, and bipartite clique". *SIAM Journal on Computing* 36(4):1025–1071.

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

LI, M. 2018. "SCOR Paper".

Naghizadeh, P., and M. Liu. 2014. "Voluntary participation in cyber-insurance markets". In *Workshop on the Economics of Information Security (WEIS)*. June 23rd-24th, State College, U.S.

N.Shetty, G.Schwarz, M.Feleghyazi, and J.Walrand. 2009. "Competitive Cyber-Insurance and Internet Security". In *Workshop on the Economics of Information Security (WEIS)*. June 24rd-25th, London, England.

Pal, R., and L. Golubchik. 2010. "Analyzing self-defense investments in internet security under cyber-insurance coverage". In *2010 IEEE 30th International Conference on Distributed Computing Systems*. June 21st-25th, Genova, Italy, 339–347.

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2014. "Will cyber-insurance improve network security? A market analysis". In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*. April 27th-May 2nd, Toronto, Canada, 235–243.

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2018. "Improving Cyber-Security via Profitable Insurance Markets". *ACM SIGMETRICS Performance Evaluation Review* 45(4):7–15.

Pal, R., Z. Huang, S. Lototsky, X. Yin, S. De, N. Bodhibrata, M. Liu, J. Crowcroft, and N. Sastry. 2021. "Will Catastrophic Aggregate Cyber-Risk Management Thrive in the IoT Age?: A Cautionary Economics Tale for (Re) Insurers and Likes". *ACM Transactions on Management Information Systems* 12(2).

Pal, R., Z. Huang, X. Yin, M. Liu, S. Lototsky, and J. Crowcroft. 2020. "Sustainable catastrophic cyber-risk management in IoT societies". In *Proceedings of the 1994 Winter Simulation Conference*, edited by S. K. S. L.-M. Z. Z. T. R. K.-H. Bae, B. Feng and R. Thiesing, 3105–3116. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Pal, R., Z. Huang, X. Yin, S. Lototsky, S. De, S. Tarkoma, M. Liu, J. Crowcroft, and N. Sastry. 2020. "Aggregate Cyber-Risk Management in the IoT Age: Cautionary Statistics for (Re) Insurers and Likes". *IEEE Internet of Things Journal* 8(9).

Pal, R., K. Psounis, J. Crowcroft, F. Kelly, P. Hui, S. Tarkoma, A. Kumar, J. Kelly, A. Chatterjee, L. Golubchik et al. 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 (TMIS)* 11(2):1–38.

Pfleeger, S. L., and R. K. Cunningham. 2010. "Why Measuring Security is Hard". In *IEEE Symposium on Security and Privacy*.

Pooser, D. M., M. J. Browne, and O. Arkhangelska. 2018. "Growth in the Perception of Cyber Risk: Evidence from US P&C Insurers". *The Geneva Papers on Risk and Insurance-Issues and Practice* 43(2):208–223.

Shetty, S., M. McShane, L. Zhang, J. P. Kesan, C. A. Kamhoua, K. Kwiat, and L. L. Njilla. 2018. "Reducing Informational Disadvantages to Improve Cyber Risk Management". *The Geneva Papers on Risk and Insurance-Issues and Practice* 43(2):224–238.

Tanczer, L., I. Steenmans, I. Brass, and M. Carr. 2018. "Networked world: Risks and opportunities in the Internet of Things".

T.H.Cormen, C.L.Leiserson, R.Rivest, and C.Stein. *An Introduction to Algorithms*. MIT Press, 2001.

Welburn, J. W., and A. Strong. 2019. *Systemic cyber risk and aggregate impacts*. RAND.

AUTHOR BIOGRAPHIES

RANJAN PAL is a faculty member of ECE at University of Michigan. His primary research interest is in engineering robust cyber-security and information privacy solutions using decision and the applied mathematical sciences. He got his PhD in Computer Science from the University of Southern California's Viterbi School of Engineering by designing pioneering market models to improve cyber-security using cyber-insurance. Ranjan did his postdoctoral research at both USC (ECE), and the University of Cambridge (CST and DPMMS). He is a member of the IEEE and the ACM, serves as an Associate Editor of the ACM Transactions on MIS, and is a network member of the Cambridge Trust and Technology Initiative. His email address is palr@umich.edu.

TAOYAN LU is an undergraduate student in ECE at the University of Michigan. His research interests include computational complexity and analytics of cyber-risk management. He is a student member of the IEEE. His email address is taoanlu@umich.edu.

PEIHAN LIU is an undergraduate student in Mathematics at the University of Michigan, specializing in data science. His research interests include computational complexity, coding theory, combinatorics, and analytics of cyber-risk management. He is a student member of the IEEE. His email address is paulliu@umich.edu.

XINLONG YIN is a graduate student in the College of Computing at the Georgia Institute of Technology. He got his undergraduate degree in Computer Engineering from the department of Electrical Engineering and Computer Science from the University of Michigan Ann Arbor. His research interests include cybersecurity-related machine learning, networks and distributed systems, and statistics. He is a student member of the IEEE and the ACM. His email address is connory@umich.edu.