JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 2024, tyae003
https://doi.org/10.1093/cybsec/tyae003

Research paper

The barriers to sustainable risk transfer in the cyber-insurance market

Henry R. K. Skeoch [1,*] and Christos Ioannidis [1,2]

[1]Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, United Kingdom
[2] Aston Business School, Aston University, Birmingham, B4 7ET, United Kingdom

*Corresponding author. Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, United Kingdom. E-mail: henry.skeoch.19@ucl.ac.uk

Received 16 August 2023; revised 18 December 2023; accepted 4 February 2024

Abstract

Efficient risk transfer is an important condition for ensuring the sustainability of a market according to the established economics literature. In an inefficient market, significant financial imbalances may develop and potentially jeopardize the solvency of some market participants. The constantly evolving nature of cyber-threats and lack of public data sharing mean that the economic conditions required for quoted cyber-insurance premiums to be considered efficient are highly unlikely to be met. This paper develops Monte Carlo simulations of an artificial cyber-insurance market and compares the efficient and inefficient outcomes based on the informational setup between the market participants. The existence of diverse loss distributions is justified by the dynamic nature of cyber-threats and the absence of any reliable and centralized incident reporting. It is shown that the limited involvement of reinsurers when loss expectations are not shared leads to increased premiums and lower overall capacity. This suggests that the sustainability of the cyber-insurance market requires both better data sharing and external sources of risk tolerant capital.

Keywords: cyber-insurance; reinsurance; Monte Carlo simulations; efficient risk transfer; cyber-threats; security economics; insurance economics

Introduction

Cyber-insurance has attracted considerable attention in the literature as a research topic and is now a significant insurance market in its own right, with $7.2bn of direct written premium in 2022 in the US domestic market alone [1] while reinsurance brokers estimate the global market may total $14bn [2]. Commercial estimates suggest that up to 45% of premium is ceded to reinsurers in the cyber-insurance market [3, 4]. Yet, the interaction between insurers and reinsurers in the cyber-insurance market has received surprisingly little attention in the academic cyber-insurance literature in comparison to industry publications [2, 5–7]. This paper aims to help partially address this gap by considering the asymmetry of information exchange and the uncertain time profile of damage revelation in relation to the cyber-insurance market and its interaction with reinsurers. It is then questioned whether reinsurers are sufficiently incentivized to participate in the cyber-insurance market on a long-term basis given the significant difficulties in achieving *ex post* efficient information exchange. Cyber risks are a relatively new multifaceted phenomena and the type of incidents and their impact may change in an unanticipated manner. It is, therefore, important to understand the resultant issues that may arise and the ability of the market to absorb unexpected losses as otherwise the sustainability of the market is threatened.

Insurance market structure

We now briefly review the structure of the insurance market and the interaction of its various associated entities and parties. A thorough analysis of the cyber-insurance market requires the role and function of the different participants in the market to be defined [1].

We assume here that the insurance buyer is a firm who buys insurance coverage via an insurance broker. The broker obtains quotes from different insurance firms provided by their underwriters. An underwriter is responsible for managing a book of insurance policies to deliver specified performance targets. These may vary according to the experience and skill of the underwriter (underwriters with a proven track record may be permitted to write either more premium

1 For further background, [8] provides a useful summary of the functioning of the Lloyd's market for insurance.

or cover riskier entities than less experienced colleagues), the markets they cover and the risk tolerance of the provider of the insurance capital. Contrary to what might be expected, underwriting is not purely a statistical exercise. The dynamics of the exchanges between underwriters and brokers are complex, in particular with respect to information exchange, which may be highly asymmetric. The job of the underwriter is to make a subjective judgement on the likelihood of the risks (prospective policyholders) they are presented with experiencing a loss and whether these can be underwritten at a premium rate, which the underwriter believes is likely to be profitable. This judgement requires a certain amount of skill as while a high insurance rate is more profitable, it will attract less demand than a more attractive rate. The key objective is to price the policy such that the desired mix of risk characteristics is obtained by the insurance firm. Underwriters are assisted by actuaries, who provide mathematical support, such as technical pricing models, to assist the underwriting-process and often hold specialist qualifications [9].

While the underwriter is the key decision maker at each insurance company in our model, insurance companies usually have multiple underwriters with different areas of expertise in terms of peril and geography—by writing policies covering different perils, insurance companies can reduce their average expected loss by diversification. Insurance brokers act as the intermediary between the insurance company and its underwriters and the end-user of the insurance. For corporate insurance, companies will typically ask their broker to prepare an insurance proposal covering a range of potential losses; these are known as lines in the insurance industry. Property, Casualty & Professional (Liability), Aerospace, and Maritime are well-known examples. The role of the broker is to obtain the best possible terms for its clients—both in terms of premium and depth of coverage. This requires the broker to have an excellent knowledge of the different insurance firms in the market and their reputation. Underwriters will aim to build a strong business relationship with leading brokers in the hope that they will receive a strong allocation of available premium.

Reinsurance companies provide insurance to insurance companies. The main reason for their existence, informally, is to smooth the potential loss profile of insurance companies who otherwise might only be able to write more modest quantities of premium or hold greater capital reserves to cover potential rare outsize losses. Reinsurers also act as a potential clearinghouse for information within the market as the reinsurer will have visbility over the portfolio contents of a range of insurers (known as cedents, which rival insurance companies in the market cannot directly observe.

Cyber-insurance presents a particularly interesting case of insurance market dynamics. The nature of the insured is particularly important as a large firm with high turnover is likely to present a more interesting and economically lucrative target for attackers, but may have better defences than a smaller firm. However, barring a systemic vulnerability the risks of significant losses in a well-diversified portfolio of numerous low-limit small-medium enterprise policy may be a more profitable undertaking for a firm. An insurance company will usually obtain reinsurance to manage either tail risks associated with its portfolio (excess of loss) or to reduce its overall exposure (quota share)[2].

Technological advancement, information deficiency, and cyber-insurance

One particular issue for understanding loss risks stemming from cyber-incidents is the difficulty in framing the potential future scope of losses. Kurz (2023) [12] provides a very readable account of the attendant challenges for economic reasoning related to advances in technology. Estimates of significant loss are usually calculated and reported by the exposure management department of an insurance company and may be either probabilistic or deterministic (based on stated realistic disaster scenarios). Exposure management traditionally is used to ascertain the risks from a natural catastrophe. In this scenario, the attacker is nature and the vectors are either wind (hurricane) or water (flooding). The questions the model for premium calculation must address are the geographical scope of the damage, which determines the expected frequency of claims and the ferocity of the natural disaster, which determines the expected severity. While nature is inherently unpredictable, nevertheless past experience of weather patterns gives some basis for modelling expected future losses. The relatively brief (at least in the history of insurance) history of cyber-risks and the constant evolution of technology, its integration in an ever increasing number of processes and the sophistication and capability of attackers makes such comparative predictions regarding potential losses extremely difficult. When designing cyber-insurance policies, it is important for the insurer to be highly specific in terms of the coverage and for the reinsurer to have a clear understanding of the risk dynamics it assumes if these policies are ceded. Figure 1 outlines a range of possible cyber losses organized by frequency and severity. This relatively simple graphic encapsulates the potential modelling challenges associated with cyber-insurance and reinsurance. The scope for cyber-losses is determined by the evolution of technology; at the time of writing, generative artificial intelligence and quantum computing are examples of two emerging technologies that have significant security implications.

What claims might arise in relation to cyber-insurance?

Barely a day passes without news of an emerging cyber-incident or other risk. It is important to realize that while these may be extremely disruptive for individuals, companies or societies, not every cyber-incident results in an insurance loss. An insurance loss may be defined as a loss resulting in a claim being paid by an insurer, whereas an economic loss is the total loss to an insured from the peril. Cyber-insurance policies comprise both first-party and third-party risks. First-party losses usually include preincident support, postincident support, cyber extortion, damage to digital assets, and business interruption, while third-party losses tend to encompass liability from privacy violations related to data loss [13]. The exact coverage will vary from carrier to carrier by policy wordings.

To consider a few different potential examples of cyber-incident losses, a ransomware attack (without data exfiltration) would largely result in first party claims for network interruption and recovery costs. In contrast, a large data breach can incur significant third party costs. It is worth noting that such third-party claims might arise several years after the policy is written. This is an important feature in cyber-insurance; a prominent relatively recent example was large hotel chain Marriott suffered a cyber-attack commencing in 2014 that was undetected until September 2018[3], which has been one of the largest cyber-insurance claims seen thus far. A reinsurer might have expected to retain the bulk of cedent premium income, only to find large claims emerging later. It should also be noted that not all insured losses related to cyber-insurance policies necessarily arise from malicious cyber-attacks. A recent example of litigation is the use of 'tracking pixels', which create a risk of claims related to privacy [15].

2 See, e.g. Albrecher (2017) [10], Kiln (2017) [11] or any other introductory resource on the fundamentals of reinsurance.

3 This data breach is widely documented on the WWW from a variety of sources; for an insurance perspective see, e.g. [14].

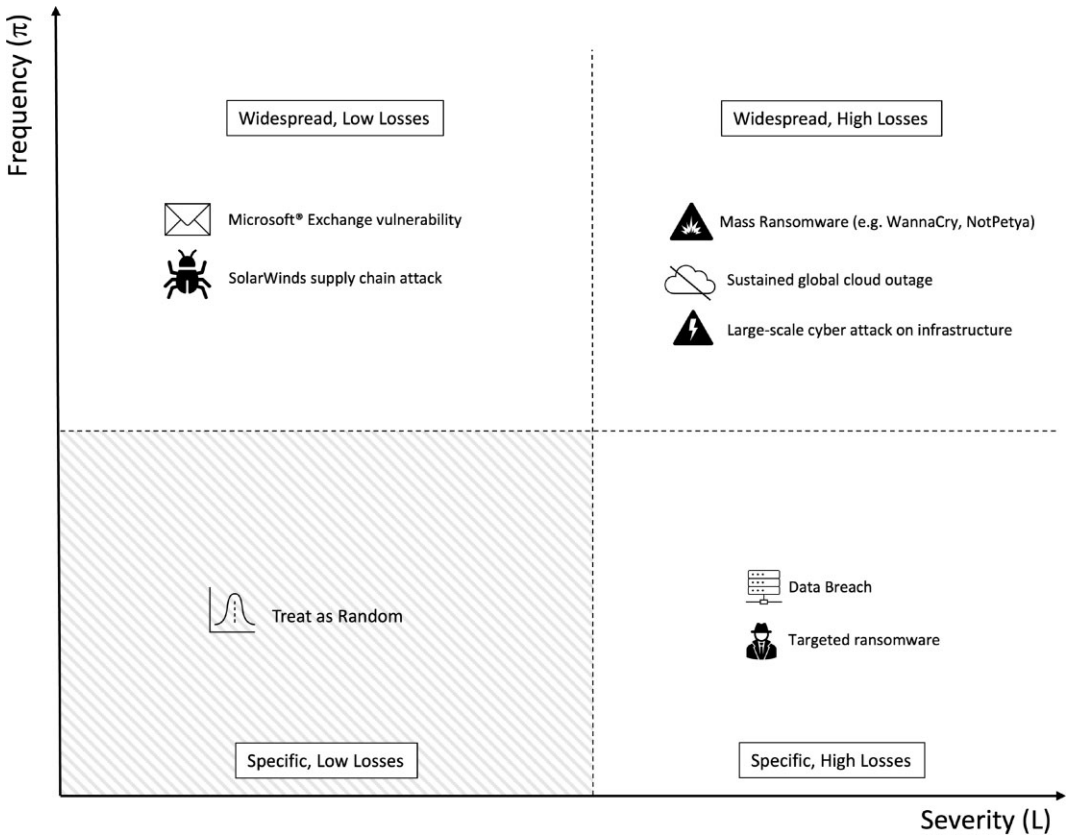

Figure 1. Categorization of cyber losses by frequency and severity.

As cyber-risk is such a nascent class of business, the insurance industry is still adapting to understanding how to price the risks, which sectors are most vulnerable and how best to assess underwriting standards. This creates a risky environment to the reinsurer, particularly as cyber is likely to be a relatively small line in their overall business and they may, therefore, lack the requisite technical expertise to truly evaluate the risks. An interesting example is Solarwinds vulnerability[4], which proved very widespread. However, it appears that the main motivation of the attackers was espionage rather than financial gain and consequently, bar investigative costs, there is little likelihood of significant cyber-claims as a result.

Figure 1 provides a stylized framework for categorizing claims related to cyber-losses. In the insurance industry, it is conventional to separate attritional losses (those experienced from ordinary course of business claims) from catastrophic losses [17]. Catastrophic cyber-losses are represented by the upper-right quadrant of Fig. 1. For the purposes of this research, we do not distinguish between attritional and catastrophic losses, but rather consider catastrophic losses as residing within the long tail (i.e. probability of loss less than 5%) of the distributions we simulate. The uncertainty for insurers and reinsures related to potential catastrophic cyber-risk is currently addressed by using third-party vendor models. A report by reinsurance broker Guy Carpenter (2023) [2] provides a useful account of the different models and their development. There remains significant divergence in the outputs of the three most commonly used models, which suggests that they are unlikely to as yet form a vehicle for agreement on the most likely distribution of cyber-losses at this juncture.

What is the motivation for reinsurance involvement in the cyber-insurance market?

A possible motivation for early entrants to the cyber-reinsurance market is to build market share and hope to capture premium rate increases as the market becomes more popular. Gallagher Re (2022) [5], a leading reinsurance broker with a specialist focus on cyber, have argued that reinsurers, technological solutions and cyber-security practice may converge to create a 'virtuous cycle of capital protection'. As insurers gain more knowledge about the likely distribution of losses, underwriting standards may be tightened. There is nevertheless a clear information advantage possessed by the ceding insurer about the 'quality' of their insurance portfolio relative to the reinsurer, which raises the issue of adverse selection. A rational reinsurer will pay extremely close attention to the information they are given by the cedent with the past loss history of the portfolio often a key feature. The fact that so much premium is ceded suggest also that insurance carriers are themselves nervous of the quantity of risks insured relative to the likelihood of losses. This begs the question as to why reinsurers would rationally increase capital allocations to the cyber-insurance market if the originating insurer is not comfortable with the risks. One possibility is that the reinsurer may have extended scope to absorb losses from cyber-risks more readily in a diversified portfolio and may further be able to charge elevated premia if the cedent insurer is desperate to offload the risk.

Industry reports suggest that there are significant capacity constraints on the appetite of reinsurers to assume cyber-risk [18–20]. A recent report by the Geneva Association[5] (2023) [7] discusses the reasons for this from an industry perspective and steps toward

4 See, e.g. Devanny *et al.* (2021) [16].

5 The leading global association of insurance companies.

improving risk sharing including broader re/insurance participation. This paper aims to formalize and model some the challenges that must be addressed to improve the efficiency of cyber-risk transfer in the insurance market. This is achieved by developing a model of a cyber-insurance market, which is stylized yet aims to be representative of the existing cyber-insurance market. The model provides a framework for mathematical evaluation of qualitative economic arguments on efficiency. A key argument of this paper is that there may be a diverse range of beliefs among market participants about the dynamics of cyber-risk and resultant losses. We demonstrate via simulations that under this assumption, reinsurance is only sometimes optimal for insurers. Economic theory on efficiency is consistent with this conclusion. This implies that insurers may need to rely on external sources of risk-tolerant capital (such as insurance-linked securities (ILS, which are sometimes called catastrophe bonds) [21–23]. Further, there is societal benefit in better information sharing on cyber-losses, which should see convergence in beliefs.

Relation to existing literature

This paper applies well-established economic theory to the cyber-insurance market in a novel manner. It is the first, to the best of our knowledge to consider the specific interaction of reinsurance capital and cyber-insurance via simulations that are representative of the existing market.

Paper structure

In this paper, we examine the impact of diverse anticipations regarding cyber-losses for firms (insurance buyers), insurers, and reinsurers in an artificial market. We compute the optimal involvement of reinsurance under different assumptions regarding the distribution of risks. The paper structure is as follows. In Literature review, we provide a brief introduction into existing literature on reinsurance, relevant actuarial models for cyber-insurance, and potential sources of external and capital requirements. Model introduces a model for interaction between the participants in the cyber-insurance market—buyers, insurers, and reinsurers. Simulations presents simulations of the artificial market under a variety of informational assumptions regarding the frequency of incidents and their severity (expected loss per attack). Finally, Discussion summarizes our results and and presents our conclusions.

Literature review

This literature review focuses on academic papers related to the fundamentals of reinsurance and on technical actuarial papers of direct relevance to cyber-insurance. It then considers the economic literature on efficiency, from which the core arguments of inefficiency in the cyber-insurance market are developed including a brief discussion of policy wordings around cyber-war. Where necessary, industry sources and commentary on the cyber-insurance market are used to illustrate arguments where peer reviewed literature is not available; but the reader should note that these sources may have commercial motivations and, thus may be less objective than comparable peer-reviewed literature.

Reinsurance fundamentals

Dionne *et al.* (2013) [24] is an excellent collection of important papers related to reinsurance and includes many key contributions to the field. Within this, Borch (1962) [25] is of particular relevance to this work, focusing on describing the conditions required to achieve equilibrium in a reinsurance market via generalizing the classical theory of commodity markets to include uncertainty. Schlesinger and Doherty (1985) [26] provide a useful treatment of issues associated with incomplete insurance markets, in particular suggesting that focusing on correlation of risks is essential for making use of incomplete markets theory. This is an argument as to why an insurer who does not currently offer cyber-insurance might enter the market should it believe that cyber-losses will not be highly correlated with areas in which it currently has exposure. Empirically, there is concern of hidden or 'silent' cyber-risks within existing lines, meaning that for many insurers offering cyber-insurance could be utility detracting. Froot and O'Connell (1999) [27] discuss the pricing of US catastrophe insurance with some illustrative data. They find that price increases and quantity declines are more pervasive than they should be within catastrophe reinsurance based on fundamental data; this is strongly suggestive of historical inefficiency.

Actuarial models

Some interesting literature has emerged around developing specific actuarial models for cyber-insurance. Bessy-Roland *et al.* (2021) [28] introduce a multivariate Hawkes process for cyber-insurance and demonstrate how it can be calibrated using the Privacy Rights Clearinghouse database of data breaches to provide a full joint distribution of future cyber attacks (see also Hillairet *et al.* (2021) [29] for an application of such modelling to cyber-insurance derivatives). Hillairet and Lopez (2021) [30] propose a stochastic diffusion model for estimating the propogation of cyber-incidents within an insurance portfolio. Biener *et al.* (2015) [31] outline a framework for systematically analysing the insurability of cyber-risk, concluding that there are significant hindrances towards a sustainable cyber-insurance market developing based on their criteria. Eling and Wirfs (2019) [32] use extreme value theory to estimate cyber-risk costs based on an operational risk database. In developing a model for cyber-insurance claims, catastrophic claims are a significant concern. Baldwin *et al.* (2017) [33] use the multivariate Hawkes process as the basis of a model for estimating contagion in cyber attacks. Bessy-Roland *et al.* (2021) [34] introduce a multivariate Hawkes framework for modelling and predicting cyber attacks frequency across firms following successful cyber-attacks against a subset of the population.

Economic theory on efficiency

We now consider how to relate well-established economic arguments on efficiency to insurance of cyber-risks. A rational buyer of insurance will likely aim to purchase a policy via a market in order to achieve a price they deem acceptable (ideally optimal). The aim of a well functioning market is to match buyers and sellers of a particular good and to establish a fair price for that good. Efficiency is often used as a measure for the efficacy of risk transfer in a market and can be defined in two ways: *ex ante* (before the transaction) or *ex post* (after the transaction). *Ex ante* efficiency requires conditions, that we shall demonstrate are extremely hard to satisfy. *Ex post* efficiency can be realized but requires an exchange of information. It should be noted that an efficient *ex post* premium if realized would be considered a technical premium (i.e. an actuarially fair premium); this is distinct and different from a premium that meets customer expectations and is subjectively viewed as acceptable based on risk tolerance or beliefs. A lack of efficiency does not mean that transactions will not take place, but creates a comparative advantage for the party with greater access to, or possession of less noisy, information.

Ex ante efficiency

According to the Arrow-Debreu model [35, 36], a complete market has:

(1) Negligible transaction costs and therefore perfect information.
(2) Every asset has a price in every possible state of the world.

Both of these assumptions are highly unlikely to be valid for cyber-insurance markets. For an asset such as a stock or bond, which may be continually traded, price is a legitimate marker of information. However, commercial insurance contracts are struck at discrete time periods and are valid for a specific length of time only. These typically operate on a yearly basis with key renewal points throughout the year dictated by market convention. Further, there is a significant cost of operating for the insurance company that is typically passed onto the customer via the premium. Most insurance contracts are nonfungible and nontransferable, unlike many publicly traded financial instruments such as stocks or bonds. In the absence of capital requirements, the purchaser of insurance is exposed to counterparty risk. This is a fundamental feature of insurance markets and implies that the first condition of completeness within the Arrow–Debreu model is unlikely to be satisfied. The second assumption that every asset has a price in every possible state of the world is equally not realistic as in reality insurance companies may decline to quote for a particular policy if the insuring party considers the risks outside of their tolerance.

Ex post efficiency

Starr (1973) [37] suggests that a set of valuation decisions is *ex post* efficient if that 'there be no redistribution that will increase some trader's realized utility while decreasing no trader's realized utility'. Alternatively, as interpreted by Feiger (1976) [38], 'there exists no alternative feasible set which is sure to be Pareto improving, looking back from the state which actually occurs.' The Arrow interpretation of states of the world is convenient for an insurance analysis as certain states of the world are loss-triggering. There are a diverse range of possibilities for attempting to frame these states—one possible utility driven approach is to model the utility of the protector of an information set using confidentiality, integrity, and availability and constructing potential incidents degrading these properties in terms of deviations from their preferred state. A cyber-insurance policy can cover a wide range of potential losses, an interesting case being costs of specialist IT consultants to help diagnosis and recovery after a data breach, for example. A data breach primarily reduces confidentiality but if the system from which the data is taken is somehow modified by a malicious actor to facilitate the theft, then it is also an attack on integrity. Recently, ransomware attacks have become a prominent cyber-threat adding a further risk of loss of availability.

A particular issue for cyber-insurance is the risk of a catastrophic cyber event. A problem with establishing distributions for catastrophic events is that the sample space is often sparse as these events tend not to occur too often. Despite computer systems and networks being societally ubiquitous in most developed countries, public data about cyber-incidents and computer mishaps of the standard required to properly price cyber-insurance contracts remains lacking. Returning to the definitions of Starr and Feiger, these require careful interpretation in the context of cyber-insurance. Consider the scenario in which an entity suffers a loss as a result of a cyber-incident, which is deemed 'with high confidence' by relevant National Cybersecurity Agencies to have been state sponsored. In an efficient market, it ought to be the case that a loss is experienced and thus constitutes a valid claim. However, many insurance policies in general contain what is known as a 'war clause' or exclusion (see, e.g. Simon (1981) [39] or Woods and Weinkle (2020) [40]). In the case of conventional property damage policies, the detailed wording of the exclusion requires careful analysis and interpretation, but conditions such as attribution and geographical scope while open to dispute are far easier to articulate in comparison with state-sanctioned if not directed cyber-operations. In the event of a significant cyber-incident, the world reaches a state whereby losses are generated. These claims ought to be paid, yet there is a clear potential for legal dispute. A prominent recent example is court cases involving (separately) Merck [41] and Mondelez [42], where both parties suffered multibillion dollar economic damages as a result of malware believed to have originated from nation state-backed entities. These claims resulted from 'all risks' property damage policies rather than standalone cyber policies [43], but nonetheless illustrate the complexity of, and attention that must be paid, to policy wordings to be clear that the policy covers what both the insurer and insured reasonably expect. Uncertainty about whether claims will be paid provides a clear path towards violation of both the Starr and Feiger conditions, meaning *ex post* efficiency is *de facto* unachievable. The state of uncertainty around 'cyber-war'[6] has led the market to take clarificatory steps. In 2022, Lloyd's of London Market Bulletin Y5381 [44] outlined requirements for state backed cyber-attack exclusions in standalone cyber-attack policies. The details of the different types of war exclusions are complex—a helpful guide is provided by the Lloyd's Market Association [45]—for the purposes of the arguments of this paper, the existence of multiple different wordings used by different organizations is supportive of our hypothesis of inefficiency in the cyber-insurance market and uncertainty around the tail risks associated with writing cyber-insurance. Wolff (2023) [46] has produced an extensive survey that relates existing literature on the role of insurance in forming *de facto* regulation to the development of war exclusions in cyber-insurance, concluding that industry leading this development may have far-reaching consequences.

Rational belief equilibria

Kurz (1994) [47] compares rational expectations equilibria, in which all agents know the true probability distribution of prices, with rational belief equilibria, in which no one knows the true distribution of prices and each agent must form their own belief about it. Even at first sight, it appears intuitive that the latter category of equilibrium is likely to better characterize cyber-insurance decisions given that a claim to know the path of future technological development with even a degree of confidence is almost certainly fallacious. Kurz's theory of rational belief equilibria relies on the system being stationary for the purposes of agents generating forecasts. The theory identifies a set $B(Q)$ of beliefs compatible with the data generated under Q, which cannot be rejected by the data. At first sight, this may appear a significant issue for analysis of cyber-risk. However, one possibility is that there exists a brittle equilibrium for a finite period of time, subject to shocks. Eventually a shock, or paradigm shift in the sense of Kuhn (1962) [48], may perturb the market from its state of equilibrium. This causes market participants to abandon their beliefs but then upon stabilization a new set of beliefs may be formed. For example, the ransomware epidemic post-WannaCry makes for an interesting case study. This introduced a hitherto less well considered generator of potential losses, which insurers had to adjust for in their policies and subsequently triggered a marked increase in premiums charged to the market based on revised distributional beliefs.

6 A full discussion of cyber-war is tangential to the work in this paper, and thus for the purposes of this research we do not distinguish between cyber-war and other catastrophic losses in the simulated market we model.

Model

We now introduce a model for describing the dynamics of a reinsurance market. We use standard results in the microeconomic theory of insurance without derivation for brevity. The motivation for this is to outline in formal economic terms the structure of an insurance market with reinsurance, from which theoretical simulations may be developed.

Insurance buyer

Before formulating the model for a market, we establish the baseline decision of a buyer of insurance facing two states—loss and no loss—with probability π and $1-\pi$, respectively. The corresponding utility function is

$$E[U] = (1-\pi)u(W - P(C)) + \pi u(W - P(C) - L + C - D), \quad (1)$$

where u is the constant absolute risk aversion (CARA) utility function,

$$u(w) = \frac{1 - e^{-\alpha w}}{\alpha}, \quad (2)$$

where α is a constant. For the purposes of this research, we chose CARA as it is a commonly used utility function and sufficiently captures the trade-offs we wish to model. Other forms of the utility function might be deployed to represent more complex buyer preferences. The parameters in Equation (1) are W, representing the initial wealth of the insurance buyer; $P(C)$, the premium paid for an amount of insurance coverage, C; and D, the deductible[7] set by the insurer. We shall assume that

$$P(C) = pC, \quad (3)$$

where p represents a premium *rate*. We emphasize that the customer chooses the coverage amount C, up to a limit permitted by the insurer and observes the premium rate, p, from different insurance companies. L is the loss experienced in the loss state. In the event that there are multiple loss states, denoted by s, we assume that these belong to a finite and countable set of states, S, such that $s \in S$, with a corresponding loss for that state, L_s. Specifying an initial endowment, W_0, and representing the total cash premium paid as P, Equation (1) may be restated

$$E[U] = \left(1 - \sum_s \pi_s\right) u(W_0 - P) + \sum_s \pi_s u(W_0 - P - L_s + C_s - D_s). \quad (4)$$

Both Equations (1) and (4) are equivalent and for the unsophisticated cyber-insurance buyer, Equation (1) is a sufficient formluation of the problem. However, when considering the supply dynamics of the cyber-insurance and reinsurance markets, it would be expected that the insurance company consider the different states that may be loss generating. We assume that the objective of the insurance buyer is to maximize their utility.

Assumption 1.
The insurance buyer aims to maximize their utility

Supply of insurance

Having established the theoretical decision framework for the insurance buyer, we now establish a formal model determining the supply of cyber-insurance. Following Hammond (1981) [49], we consider the actions of consumers in the economy:

$$x^i(s) = \left[x_t^i, x_{t+1}^i(s)\right], \quad (5)$$

where i represents an individual consumer of a total I consumers in the marketplace. As before, s represents a contingent state of the world, and it is assumed that the set of possible states, S is finite. The vector of total insurance demand, $\mathbf{x}_t = [C_t^1, C_t^2, \ldots, C_t^i]$.

We assume that there are J insurers in the market, each with a supply of insurance

$$\mathbf{y}^j(s, \mathbf{x}) = \left[\mathbf{y}_t^j(\mathbf{x}_t), \mathbf{y}_{t+1}^j(s, \mathbf{x}_{t+1})\right], \quad (6)$$

where $\mathbf{y}_t^j$ is an i-length vector of the units of insurance sold by insurer j to customer i at time t and consequently, which, expressed in monetary terms is identical to cover, C. It is assumed that each customer i has an exclusive policy with its chosen insurer j. Each insurer has a premium vector,

$$\boldsymbol{P}^j = [P_1, P_2, \ldots, P_i], \quad (7)$$

representing the premium it charges to each customer. This vector may be time dependent. For conciseness of presentation, we will henceforth drop time subscripts as the analysis in this paper is restricted to a single period.

Insurer objectives

We assume the insurer formulates its decisions on insurance supply, $\mathbf{y}^j(s, \mathbf{x})$ via the following parameters (see Chapter 3.5 of [50]):

- K: the reserve capital held by each insurer.
- P: the total premium income for each insurer.
- X: the claim costs (losses) experienced, described by probability function $F(X)$ with differentiable density $f(X)$ defined over the interval $[0, X_{\max}]$.
- D: the total deductible enforced by the insurer.
- W_0: the initial wealth of the insurer—this may be thought of as shareholder equity, e.g. or syndicate (nonregulatory) capital.
- W: the residual wealth the insurer has after paying claims. If the amount of claims is greater than $A \equiv P + K + W_0$, the insurer faces ruin.
- r: the risk-free interest rate for the relevant period.

We assume that each insurer has zero utility condition and its objective is to maximize its wealth

$$W_j = W_0 + P_j - \int_0^{A_j} \frac{X_j - D_j}{1 + r} dF_j(X) \quad (8)$$

subject to the constraint

$$W_j + K_j > 0. \quad (9)$$

The intuition underpinning Equation (8) is that the insurer has a trade-off between the amount of premium it collects and the risk of claim associated with that premium. It may also set a deductible to mitigate moral hazard. Accordingly, the insurer should assess the probable maximum loss of claims according to its distribution and ensure that it has sufficient capital to pay the claims. The claims are discounted by the risk-free rate, r, as it is assumed that the insurer will earn interest on its earned premium over the period. The optimal set of allocations for the insurer would be to policies that maximize the wealth/capital ratio W_j/K_j.

Assumption 2.
Insurers aim to maximize their wealth

7 The amount of losses, which must be borne by the insurance buyer.

Assumption 3.
The probability distribution, $F_j(X)$ is subjective to each insurer in the sense of de Finetti (1974) [51]. This will be justified in Modelling cyber-risks.

Introducing reinsurance

In order to reduce risk exposure, the insurer may also seek to purchase reinsurance. There are two categories of reinsurance considered in this work: quota share and excess-of-loss. Reinsurers are consequently concerned with determining the probability of two types of extreme events: those resulting in single large losses from a particular client (concentrated losses/attritional) and those resulting in widespread repeated claims across cedents (contagion/catastrophic). In the event that this distribution is objective, then this would lead to a universal fair price for insurance. Reinsurers in turn will have their own subjective distributions and charge the expected value of their own distributions to clients. While this may be commercially reasonable, such prices are not fair in a strict economic sense. The existence of reinsurance serves to allow insurers to smooth their subjective expected loss distributions, which clearly implies risk aversion as opposed to neutrality. In short, intermediation implies imperfection[8]. Including reinsurance, Equation (8) becomes:

$$W_j = W_0 + (P_j - R_j) - \int_0^{A_j} \frac{X_j - D_j - I_j}{1+r} dF_j(X), \tag{10}$$

where the parameters are as above, with the addition of R, which represents the cost of reinsurance to the insurer and I_j, which is the amount of losses indemnified by the reinsurance policy purchased. The constraint $W_j + K_j > 0$ continues to apply. Notation-wise, in similar fashion to Supply of insurance, we use vectors to describe reinsurance supply. We assume that there are k reinsurers, who charge $\mathbf{r}^k$ rates to insurer j and denote the supply vector of reinsurance as $\mathbf{z}^k$.

For a simple quota share policy,

$$R = \rho P,$$

where ρ is the proportion of the portfolio ceded and then

$$I(L) = \rho L.$$

However, in cases involving excess of loss or other reinsurance treaties, the calculation is more involved. Miccolis (1977) [53] provides an exposition of some standard mathematical techniques for describing excess of loss calculations. In the case of excess of loss, the indemnification equation becomes:

$$I(L) = (L - N)^+ - (L - N - M)^+, \tag{11}$$

where M and N are parameters for an excess of loss policy covering \$M(mn) of losses in excess of \$N(mn). For simplicity, it is assumed that each insurer can purchase only a single excess-of-loss policy from each reinsurer. It would seem rational for the purposes of our discussion that the insurers seek to buy reinsurance above the aforementioned value A, losses above which the firm becomes insolvent.

The reinsurance market

We assume that there are K reinsurers in the market who provide reinsurance capacity. The reinsurer aims to maximize wealth in similar fashion to the insurer (Equation 8), but does not include a deductible:

$$W_k = W_0 + R_k - \int_0^{A_k} \frac{I_k(X)}{1+r} dF_k(X), \tag{12}$$

where R_k is the total reinsurance premia received and $I_k(X)$ denotes expected reimbursements paid out to cedents. The reinsurer is subject to the capital constraint $W_k + K_k > 0$. Note that we allow for the reinsurer and insurer to have different beliefs about the expected distribution of losses. As with the insurers, A_k, represents the amount of reinsurance payouts above which the reinsurer would be insolvent.

Assumption 4.
The reinsurer may have a different belief from the insurer regarding the distribution of risks.

Modelling cyber-risks

We have, thus far considered losses related to cyber-risk in an abstract sense as setting up the theoretical framework for evaluating the interaction between buyers, insurers, and reinsurers does not require the functions dictating these losses to be instantiated. However, simulating the decision making to analyse the potential for efficiency in the market does require some sample distributions. We use standard results in probability theory without derivation (the reader wishing to understand the background more thoroughly is referred to any standard statistical text on probability theory; Williams (1991) [54] is a particularly accessible and carefully explained introduction). While using formal probability theory is not essential for simulating the results in this paper, it is beneficial to apply theoretical rigour as this helps to highlight some features specific to cyber-risk that are potentially problematic for formulating traditional actuarial insurance assessments.

We start by defining a probability triple $(\Omega, \mathcal{F}, \mathbf{P})$. Ω is a set representing the sample space of *all events*. ω represents a sample point of the sample space. The σ-algebra[9], $\mathcal{F}$, on Ω, is known as the family of events[10]. Denoting an event by A, we may write

$$\mathcal{F} = \{A | A \subseteq \Omega, A \in \mathcal{F}\}. \tag{13}$$

The intuitive explanation in relation to cyber-insurance is that $\mathcal{F}$ is the collection of events covered by a policy that may trigger a claim and then, possibly, a loss to the insurer. If $\mathcal{F}$ is the Borel[11] σ algebra on the set of real numbers, then there exists a unique probability measure on $\mathcal{F}$ for any cumulative distribution function. Letting X be a random variable on $(\Omega, \mathcal{F}, \mathbf{P})$,

$$\begin{array}{ccccc} \Omega & \xrightarrow{X} & \mathbb{R} \\ [0,1] & \xleftarrow{\mathbf{P}} & \mathcal{F} & \xleftarrow{X^{-1}} & \mathcal{B}. \end{array} \tag{14}$$

Informally, this means that so long as there is a collection of events that obeys certain mathematical properties, it is possible to assign a probability to an event using a probability distribution function. One interesting outcome is that a key assumption of probability theory is that the system is stable. This is a potentially problematic assumption for cyber-risk as there have been clear examples of previously unconsidered threats developing. However, insurance policies comprise a set of event definitions as part of the policy, which are contractually binding (albeit open to legal dispute). The importance of careful policy wording is consequently readily apparent. As will shortly be explained, underwriting cyber-insurance policies requires an assumption of subjective, temporary stationarity in distributions. This

8 Skiadas (2013) [52] presents an interesting analysis on this topic.

9 The definition of a σ-algebra is a collection of subsets of a set that is closed (stable) under any countable number of set operations. This is important for working with probabilities, where the probabilities of all possible outcomes must sum to 1.

10 See Chapter 2 of Williams (1991) [54].

11 The Borel σ-algebra, $\mathcal{B}(\mathbb{R})$, is the smallest σ-algebra containing all open intervals in $\mathbb{R}$.

is a realistic assumption in the context of industry practice, where (re)insurance policies last for a year and then are repriced based on updated distributions resulting from supply–demand dynamics and claims experienced.

Why use subjective probabilities to model cyber-risks

Assumption 3 in Supply of insurance stated that the probability distributions that govern insurance supply are subjective in the sense of de Finetti (1974) [51]. We now provide the intuition behind and justification for this assumption before moving to consider the form of distribution that might be used to model cyber-insurance decisions.

In Ex post efficiency, we outlined the conditions required for *ex post* efficiency. Considering these in the context of cyber-insurance, we conclude that *ex post* efficiency is unlikely to hold and almost certainly cannot be implemented at the time when the underwriting decision is made. Unless of course, the true probability distribution attached to the known and finite states of nature is known and shared by all participants. Such condition is the foundation of the theory of rational expectations. This is synonymous with the existence of a stationary distribution. One way of defining a stationary process is to say that its moments are time-independent, which means that the average value of the measurements is a constant. Such distributions are foundational for the existence of efficient equilibria under risk.

It is usual in macroeconomics to depict technological progress as a Markov chain. If the depicted process has started far from its invariant distribution, then it is also nonstationary, but easy to predict as it will approach the limiting distribution that is ultimately stationary. However, in a short epoch, it will appear as nonstationary. Whether technological progress has such a limiting distribution is an unresolved question. Over the long-run it appears to have exhibited a definite trend, with some downwards transitions attributable to natural disasters, wars epidemics etc. In the short run, local approximations can be derived, and expectations can be formed, however, agents will splice different segments depending upon their horizons and discount rates. The imposition of rational expectations restrictions upon this structure can only be justified if all agents have identical preferences and endowments, a condition that by construction does not hold. For Markov chains with nonstationary transition probabilities, no steady-state typically exists and almost nothing in the nonstationary setting is computable in closed-form.

It is hard to imagine that there is any way to truly predict an arbitrary nonstationary process. This is because as soon as one postulates a future path another can always reverse it, without creating any problems of consistency with earlier data. In a more general case one might lower expectations, not to actually predicting well, but to predicting with low regret. To this effect agents can choose their most suitable approximations selecting the time span and use their best computational algorithms.

In the absence of a universally accepted probability distribution, *ex post* efficiency is almost impossible to attain. Of course, there are opportunities which can best utilized only with *ex ante* knowledge of the state of nature. In its simplest form, it is the choice of technique in production/product that depends upon the expected state. However, a more interesting situation arises when the expected state conditions the preferred level of production. Expecting the cyber-insurance market to quote premia at all levels that are consistent with *ex post* efficiency is rather unrealistic. The very nature of the underlying processes does favour the existence of a generally accepted stationary distribution. Rational agents will behave as if they are *ex ante* efficient using their own expectations of losses based on their subjective probability distributions taken over their own sample spaces.

The evolution of cyber-threats will be conditioned of the path of technological improvements in both elements of information and communications technology, software and hardware. The future path of such advances may be partly predictable based on well-established empirical regularities, such as Moore's law that famously predicted that the number of transistors on integrated circuits would double every two years, i.e. at an annual rate of about 40% [55]. Others[12], looking at related data came to the conclusion that predictions of particular technological IT innovations, such as hard drives may be approximated using exponential functions. A very useful exposition of this attempt, using smooth functions the predict technological progress is Farmer and Lafond (2016) [59].

Yet, technological advances undergo structural breaks, where both the level of technology in terms of some of its main characteristics and its future direction change. A prominent example at present is the introduction of quantum computing, which will alter radically reduce computational time and, thus has implications for the robustness of cryptographic protocols that are currently infeasible to attack on a realistic timeframe.

Technological progress is achieved by the complex interactions of two main human pursuits. The organized knowledge as it appears in scientific papers, submitted patents, recipes, protocols, routines, and probably informal know-how, acquired through 'learning by doing' in a long process of imitation and repetition. The development of science, technology, innovation and production require both codification and knowledge.

It seems unlikely that such dual processes can be tamed into a smooth parametric function with time invariant parameters, shared by all participants. If anything, in the absence of such shared beliefs, it is expected that for market participants whose welfare depends upon such developments, their decision making will be based on arbitrarily diverse anticipations. These are individually efficient decision makers because they act on the basis of all the information available at the time. It is clear, therefore, that by and large insurance contracts on expected losses based of future technological developments, that are subject to structural changes, cannot be written on generally accepted parameters, to deliver Arrow–Debreu type *ex ante* efficient premia. All the participants are efficient in terms of fully exploiting their private anticipations of losses, but the quoted premia at the two levels will not result in fully efficient in the Pareto sense economic outcomes.

Probability distributions

For the Simulations, we separate the expected distribution of losses into the number of expected claims (frequency) and the average expected loss per claim (severity). This is a very common method for actuarial modeling and is described in most standard texts, for example Panjer (2006) [60]. Its appropriateness to categorizing cyber-risks was described in Section I.3 and summarized in Fig. 1. We assume that frequencies follow a Poisson distribution and severities a log-normal distribution. The Poisson distribution is a standard starting point for frequency modelling (see, among many, [61]). There is no clear consensus in the empirical literature on which distribution is most appropriate for describing the severity of cyber-losses (see in particular [32, 62]). We use the log-normal distribution as a starting point as it is well-understood and straightforward to configure. We use simulated rather than empirical distributions as the aim of the simulations is to examine whether efficiency is theoretically possible, whereas markets in practice are very unlikely to be efficient. The

12 For example, Benson and Magee (2015) [56], Funk and Magee (2015) [57], and Nagy *et al.* (2013) [58].

probability of k events occurring in a unit of time represented by the Poisson distribution is

$$f(k,\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}, \tag{15}$$

where λ is the expected number of events. The log-normal distribution assumes

$$\ln(X) \sim \mathcal{N}(\mu, \sigma), \tag{16}$$

that is the natural logarithm of variable X is normally distributed with mean, μ, and standard deviation, σ, which are defined as

$$\mu = \frac{\mu_X^2}{\sqrt{\mu_X^2 + \sigma_X^2}} \quad \text{and} \quad \sigma^2 = \ln\left(1 + \frac{\sigma_X^2}{\mu_X^2}\right), \tag{17}$$

where μ_X and σ_X^2 are the mean and variance, respectively, of the variable X. The probability density functions and cumulative distributions functions for the log-normal distribution are readily available in any standard resource on statistics and are omitted for brevity.

Combining probability distributions

Cyber-insurance policies cover a diverse range of first- and third-party risks and consequently, there is probably no one distribution that actually covers all relevant risks. Accordingly, it is desirable to consider a combination of possible risks. Unfortunately, probability distribution functions are rather difficult to combine with a closed-form solution (see, e.g. Nadarajah *et al.* (2018) [63]) and require analytical solutions. A common strategy is to use a package such as Mathematica [64]. However, there is an alternative approach which is to use Monte Carlo-type simulations. Simulations will show how these can be deployed to yield useful insights on insurance decisions, the results of which do not require sophisticated mathematics to formulate or interpret.

Simulations

We consider simulations of a cyber-insurance market with reinsurance over a single period. We assume that losses arise in the period of the insurance policy and are recorded at the time they arise. Policy data is confidential to insurance companies and consequently, the simulations are established for model convenience but are constructed to replicate real-world insurance market dynamics. We use Poisson distributions for the frequency of losses and log-normal distributions for the severity of losses (details of these distributions and their associated functions may be found in any standard statistical text). The Poisson distribution is a common choice for modelling claim frequencies in insurance (see, among many excellent references, [61, 65]). There is no clear consensus in the literature on the optimal distribution for modelling the severity of cyber-related claims, but the log-normal distribution has been shown to be a reasonable approximation in the limited empirical studies to date (e.g. Eling *et al.* (2019) [32], Woods *et al.* (2021) [62]. The use of the joint frequency-severity distribution approach follows Panjer (2006) [60]. We assume a common set of contracts across insurers varying in limit size.

The analysis considers only variation in coverage and premium. We assume arbitrarily a market size of $500mn total coverage. The simulations were computed using the Julia programming language. We found the *Distributions.jl* [66], *QuadGK* [67], and *Plots.jl* [68] packages particularly useful in facilitating the presentation analysis. Unless otherwise specified, *Monte Carlo* type simulations were run 100 000 times.

The goal of the simulations is to illustrate how capital supply from the reinsurance market to the insurance market and then to buyers is inherently inefficient as pricing is influenced by the diversity of opinions regarding the frequency and severity of losses even with relatively simple standard distributions. The simulations might be applied to a variety of insurance markets, but they have been constructed to be representative of the existing cyber-insurance market based on the authors' interaction with insurance market professionals.

Preliminaries

Familiarity with the insurance market is not a prerequisite for understanding and interpreting the simulations that follow. We have taken care to explain the terms used and ensure parameters are fully defined and explained. However, the reader unfamiliar with corporate insurance may find the following definitions helpful as a reference. These may be safely skipped for those experienced in either the practice or study of insurance.

- μ_L: the average expected loss in monetary (cash) terms.
- σ_L: the standard deviation of losses in monetary (cash terms).
- $F^{-1}(p)$: the loss value that occurs with probability p according to the cumulative distribution function F. If $P = .95$, then in 95% of cases, the loss is expected to be less than or equal to the output of this function.
- Loss ratio: the percentage of cash premiums collected by an insurance company for a specified period (usually a year) paid out as losses.
- Frequency: the number of claims in a period.
- Severity: the average loss per claim.
- Cover/exposure: the total maximum losses that could result from a policy/portfolio, respectively.
- Expected loss: the mean loss from a policy/portfolio.
- Net loss: the loss to the insurer after applying purchased reinsurance.
- Technical premium: the cash premium or premium rate (calculated as the ratio expected loss/cover) corresponding to the expected loss. This is the premium income at which the insurer can be expected to break even.
- Simulated loss: the average loss from running N simulations based on random sampling of the expected loss distribution. This can only be computed once the portfolio is formed, so we assume that premiums are calculated based on expected loss values.
- Ceding commission: the percentage premium paid back to a cedent by a reinsurer to cover underwriting expenses and other costs.

It is important to note the sequencing of the insurance transactions in the market. The insurance buyers observe a premium rate and based on this decide how much cover to buy. The insurance provider then has obtained a portfolio. Based on the risk characteristics of that portfolio, the insurer may look to enter into a reinsurance contract to eliminate some potential risk. The simulations assume that insurers and reinsurers target a specific loss ratio *ex ante* to determine pricing.

Simulation strategy

We consider three simulations:

(1) A benchmark simulation.
(2) One reinsurer, five insurers with different portfolios comprised of different weights of five common contracts, buyers not considered.
(3) One insurer, one reinsurer, different buyer price sensitivities.

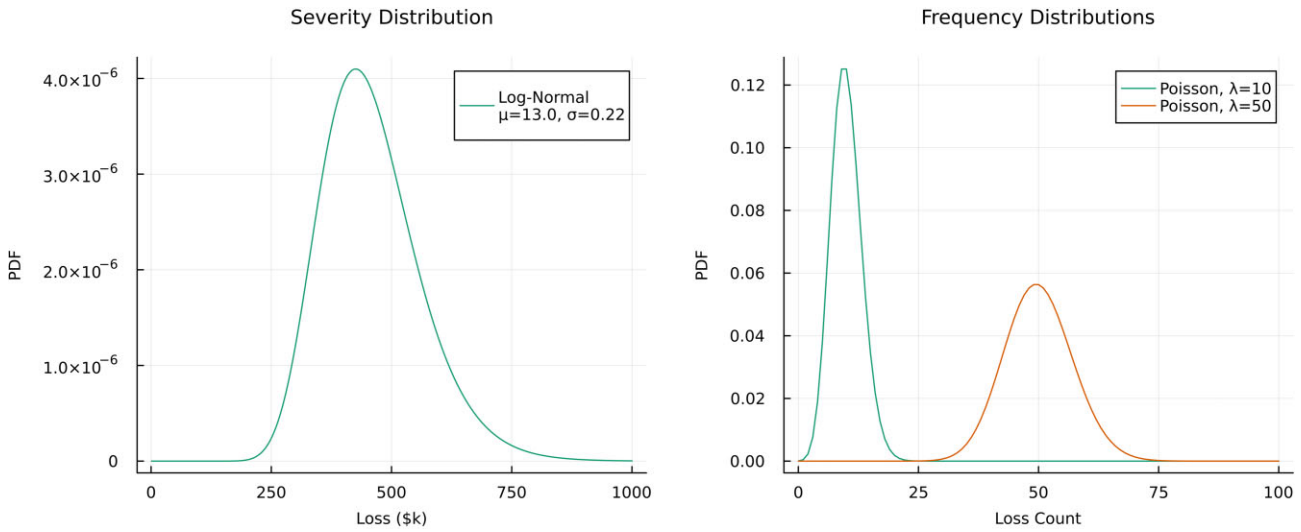

Figure 2. Benchmark severity (LHS) and frequency (RHS) distributions.

These simulations are distinct from each other, though have broadly consistent parameters where possible. The aim of the benchmark simulation is to demonstrate the approach used to generate loss distributions and also to instantiate buyer utility functions to show that if the buyer has a different expectation of loss severity from the insurer, then full insurance coverage may not be utility maximizing.

The second simulation starts with a reinsurer who has a range of distributions its actuaries consider acceptable. The reinsurer attempts to offer reinsurance to achieve a target loss ratio and so quotes a reinsurance rate to the market. The market consists of five insurers who have portfolios that range from a large number of small loss risks (called Insurer Alpha) to a small number of large loss risks (called Insurer Echo) with Insurers Beta, Charlie, and Delta having portfolios that move progressively between the two extremes. This aims to replicate the structure of the cyber-insurance market in a stylized form and contrast the appropriate reinsurance strategy for the different types of insurer.

It should be noted that the premia in the simulations may vary from those witnessed in the market and in some cases appear very large. The simulations are intended to guide the reader through an application of the economic theory and market model from a theoretical perspective and demonstrate the difficulty of establishing efficiency rather than aiming to be a simulation of the real-world cyber-insurance market.

Simulation 1: benchmark simulation

We first consider a simple simulation before starting to examine the effects of varying market structure and pricing variables. This simulation assumes the following:

- There is only one insurance policy offered in the market, with a limit of $1mn.
- The mean expected severity of an incident (loss) for each policy is $500k, with standard deviation $250k.
- The frequency of losses is simulated under two scenarios where 10% and 50%of policies are expected to experience a loss, respectively.
- There are 100 buyers, five insurers, and one reinsurer in the market. For simplicity, we model total losses for the market and assume they are evenly distributed.
- Losses are simulated with 100 000 runs and random sampling of the severity and frequency distributions.
- Distributions are shared by all market participants.

Figure 2 plots the probability distribution functions of the severity distribution and the two frequency distributions. The severity distribution is log-normal with parameters $\mu = 13.0$ and $\sigma = 0.22$; the two frequency distributions are Poisson with λ of 10 and 50, respectively. The PDF values for the severity distribution are very small because of the units of the loss; the integral of the PDF across the function domain must sum to 1. Running a simulation, the expected loss distribution for the two frequency distributions can be obtained. This is presented in Fig. 3. The values on the *y*-axis of Fig. 3 simply represent the number of times each loss value range in the histogram appears in the simulation. Each bar in the histogram has a width of $0.5mn. This is simply chosen for aesthetic reasons. The main emphasis is on the shape of the distributions rather than the precise frequency count in the histogram.

Having examined the distributions, we now consider the pricing of the policies. Table 1 shows the expected and simulated losses for the distributions in Fig. 3. Note that

$$\text{Expected Loss} = \text{Expected Frequency} \times \text{Expected Severity} \times \text{Number of polices.}$$

The ratio of the expected loss and the exposure ($100mn in this example) gives what is known in insurance as the technical premium rate. Accordingly, the technical premium would be 5% for the 10% frequency scenario and 25% for the 50% frequency scenario. The simulated losses are lower than the expected (mean) losses because of the skew of the log-normal distribution.

To simulate reinsurance pricing, we first fit a log-normal distribution to the joint distribution with 50% loss frequency as previously described. This should be understood as the reinsurer attempting to estimate the 'true' underlying distribution and is an approximation. We consider reinsurance only for the 50% loss frequency distribution as guided by the reported loss ratios in Table 17, which suggest relatively high frequencies of losses have been experienced by the actual market. Using the *fit* functions in *Distributions.jl*, we obtain a log-normal distribution with $\mu = 16.9$ and $\sigma = 0.27$. Under this distribution, the cumulative probability of a loss exceeding $50mn

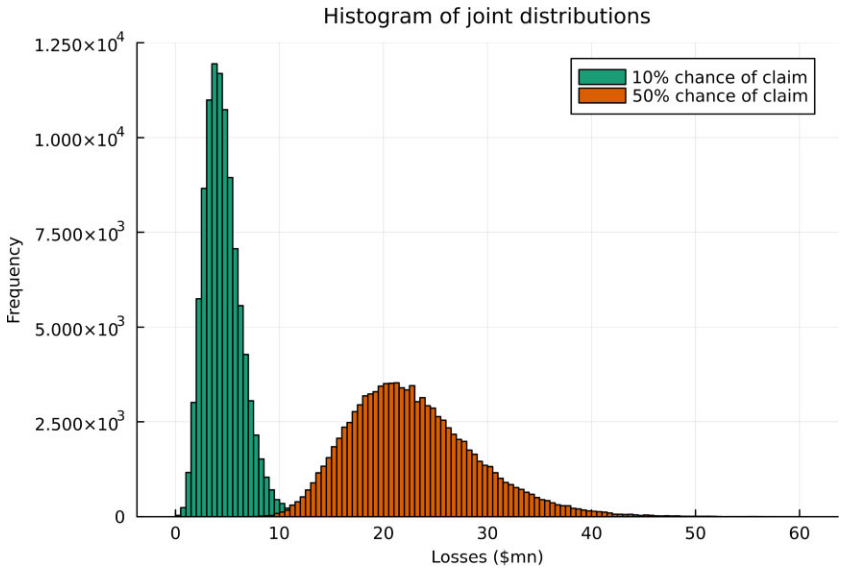

Figure 3. Simulated loss distributions.

Table 1. Expected versus simulated benchmark distribution losses.

Frequency (%)	Expected loss (mn)	Simulated loss (mn)
10	$5	$4.6
50	$25	$22.9

is extremely small, therefore, we price the reinsurance policies using excess-of-loss policies with different attachment points and limits up to an exhaustion point of $50mn. Using the cumulative probability functions for the estimated distribution, we can then obtain premium rates for the reinsurance, which, multiplied by the amount of reinsurance required, gives the cost of reinsurance. We then rerun the simulations of losses for the insurer assuming no losses are incurred above the threshold at which reinsurance cover binds. We can then obtain the simulated loss with reinsurance. The results are presented in Table 2.

This reinsurance pricing may be considered efficient because both the reinsurer providing coverage and the insurer seeking reinsurance have the same expected loss distribution. It should be noted, however, that the difference between the simulated net loss and the $22.9bn in Table 2 exceeds the technical reinsurance premium. In theory, the two should be equal. The discrepancy arises as a result of fitting error in approximating the joint Poisson/log-normal distribution via a log-normal distribution. This first simulation is intended only as an overview of how to price reinsurance, and so this discrepancy is noted simply for transparency and to avoid potential confusion. It has no impact on the simulations in the following sections where the reinsurer and insurers do not share distributions.

Simulation 2: reinsurance supply and price

Having considered the case where all parties agree on the same distribution, we relax this assumption and start to consider divergence in distributions of expected losses (Assumption 4 in Introducing reinsurance). The justification for this is the heterogeneity of loss ratios in the US cyber-insurance market in Table 17, which are discussed in detail in Cyber-insurer loss experience. We begin by considering the objective of the reinsurer. We assume a log-normal distribution of total losses. This is the distribution the reinsurance company believes represents the losses experienced from a pool of cedents. The reinsurance company needs to model different potential loss ratios. Initially, we assume cover is fixed at a maximum of $500mn. Table 3 presents a number of log-normal distributions. These are purely for illustrative purposes; in a real world situation, the reinsurer would model the distribution based on experience and data. However, it is helpful to consider a range of distributions to understand how the shape of the distribution may affect pricing.

Within this table, $F^{-1}(0.995)$ represents the maximum loss with 99.5% certainty within the distribution. This is the probability value used under the Solvency II insurance regulation to determine the required capital a firm must hold. The probability density function and cumulative distribution functions for the distributions in Table 3 are plotted in Fig. 4. Note that the scale of the loss axis is shortened to $100mn as the probability density function returns extremely low values beyond this point.

To estimate the premium rate, we consider the following. The reinsurer targets a loss ratio (a common performance metric in the insurance industry). Total losses from the portfolio are then written

$$L = \text{L.R.} \times \sum_J r_J C_J. \tag{18}$$

Losses experienced are also given by

$$L = \sum_J E[I_J]. \tag{19}$$

We assume there is a single rate for reinsurance such that $r_J = r \forall J$. Then,

$$r = \frac{\sum_J E[I_J]}{\text{L.R.} \times \sum_J C_J}. \tag{20}$$

Denoting $\bar{C} = \sum_j C_J$ and noting that $\sum_J E[I_J] = \int_0^{\bar{C}} I f(I) dI$ where $f(I)$ is the probability density function of an appropriate distribution, we obtain

$$r = \frac{\int_0^{\bar{C}} I f(I) dI}{\text{L.R.} \times \bar{C}}. \tag{21}$$

This integral can be evaluated numerically, for example using *QuadGK* in Julia.

Table 2. Reinsurance premia for excess-of-loss policies on the 50% chance of claim distribution.

Reinsurance	Reinsurance premium rate (%)	Reinsurance cover (mn)	Technical reinsurance premium (mn)	Simulated net loss (mn)
$25mn xs $25mn	32.2	$25	$8.1	$13.1
$20mn xs $30mn	12.5	$20	$3.8	$18.7
$15mn xs $35mn	4.2	$15	$0.6	$21.3
$10mn xs $40mn	1.3	$10	$0.1	$22.4

Table 3. Table of reinsurance distributions.

		μ_L (mn)	σ_L (mn)	μ	σ	$F^{-1}(0.995)$ (mn)
Distributions	A	$10	$10	15.8	0.69	$42
	B	$20	$20	16.5	0.69	$84
	C	$30	$30	16.9	0.69	$126
	D	$40	$40	17.2	0.69	$169
	E	$50	$50	17.4	0.69	$211
	F	$60	$60	17.6	0.69	$253

Suppose the reinsurer believes that Distribution C best describes expected losses to the portfolio and targets a loss ratio of 50%. The rate of reinsurance charged is then 11% (Table 4). Premium income for the reinsurer will be $55mn. Note that per Table 3, in Distribution C, the 99.5% upper bound for losses is $126mn. This means that the reinsurer must have access to an additional $71mn of capital under this policy scenario.

Insurance supply

We assume for simplicity that there are five insurance contracts in the market with different limits: $500k, $1mn, $2mn, $5mn, and $10mn. We assume that there is a uniform individual and independently distributed probability of loss for each contract: The expected severity in the above contracts is assumed to be log-normally distributed per Table 5 and the frequency $\sim$Poisson($\pi_l k$) where k is the number of contracts. Table 6 contains a sample portfolio for a panel of five insurers for illustrative purposes to run a loss simulation. The technical premium is the premium income that equates to the expected loss. Equivalently, this is the premium written at which the insurer would expect to break even.

In reality, insurers do not attempt to break even but rather aim to produce a profit to provide a return on investment to the source of their capital. One simple objective might to not exceed a target loss ratio. This is achieved via an additional charge to the insurance buyer over the technical premium known as a loading[13]. The loading is calculated:

$$\text{Loading} = \frac{1}{\text{Total Exposure}} \times \left(\frac{\text{Technical Premium}}{\text{Target Loss Ratio}} - \text{Technical Premium} \right). \quad (22)$$

The variation between loading and loss ratio for the insurance portfolios in Table 6 is plotted in Fig. 5. The variation in target loss ratios may occur for a number of reasons, such as rate of return on capital demanded by the capital source (as discussed in Insurance market structure, prior loss experience, or other variable expenses. The loading also may aim to capture any skew in the actuarial distribution.

Table 7 shows the calculated loadings for each insurer in the simulation assuming a target loss ratio of 50%. For ease of comparison, we keep the target loss ratio constant across the insurer panel and also the overall exposure.

Interaction between insurance and reinsurance

The total expected losses for the cyber-insurance market depicted in Table 7 are $38.0. The technical premium is equal to the expected losses in monetary terms. In Simulation 2: reinsurance supply and price we stated that for a log-normal distribution with mean and standard deviation of $40mn and target loss ratio of 0.5, the premium charge would be 14% for the reinsurer (distribution D). The usual process of reinsurance in quota share is that the reinsurer assumes a stated percentage of portfolio losses. The reinsurance contract (or treaty) is priced[14] via a ceding commission and reinsurance margin. In this case, the reinsurance margin is already accounted for in the 14% premium rate as this was calculated to give the required reinsurer loss ratio. The ceding commission is paid back to the ceding insurer to compensate them for underwriting expenses. The ceding commission is defined as the average premium rate (Table 7) less the cost of reinsurance (14%). Inspecting Table 7 once more, we can see for insurers Charlie, Delta, and Echo, the average premium rate of the portfolio exceeds the reinsurance cost. Therefore, the ceding commission for these insurers would be positive. However, for insurers Alpha and Beta, their weighted average premium rate is below that charge for reinsurance, implying a negative ceding commission. If Alpha or Beta believe that their assumed distributions are correct, this would not be rational behaviour. For the other insurers, purchasing reinsurance would reduce profits for *expected losses*. However, the value of reinsurance will become apparent once we consider the effect of capital.

Having established the target pricing for each insurer *ex ante*, we now consider simulating *ex post* losses. The profit equation for the insurer, may be written:

$$\begin{aligned} \$\text{Profit(L)} = {} & \$\text{Premium Written} \times (1-\rho) \\ & + \$(\text{Exposure} \times \rho \times \%\text{Ceding Commission}) \\ & - L, \end{aligned} \quad (23)$$

$$L = \begin{cases} \text{Loss}(1-\rho) & \text{if } D = 0 \\ \text{Loss} & \text{if } D > 0, \text{Loss} \le D \\ D + (1-\rho)(\text{Loss} - D) & \text{if } D > 0, \text{Loss} > D \end{cases}, \quad (24)$$

where ρ is the fraction of the portfolio ceded to the reinsurer and D is a deductible. We restrict our analysis in this simulation solely to policies without deductibles, but provide for their inclusion for completeness.

Simulation procedure

For each insurance portfolio in Table 5 we simulate losses via the following procedure.

13 See, e.g. Benjamin (1986) [69] for a discussion.

14 Clark (2014) [70] is a highly approachable introducing to reinsurance pricing.

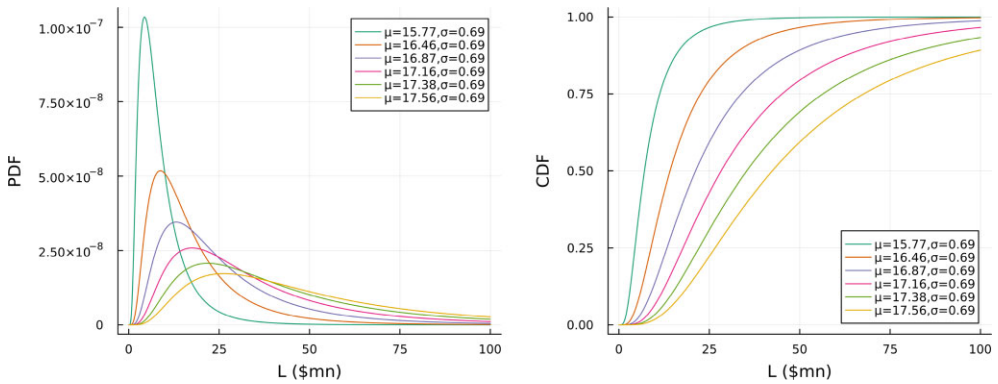

Figure 4. Reinsurer loss distributions.

Table 4. Illustrative premium rates for target loss ratios under different distributions at cover fixed at $500mn.

		Premium rates								
Loss ratios→		0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Distributions	A	0.18	0.09	0.06	0.04	0.04	0.03	0.03	0.02	0.02
	B	0.36	0.18	0.12	0.09	0.07	0.06	0.05	0.04	0.04
	C	0.54	0.27	0.18	0.13	0.11	0.09	0.08	0.07	0.06
	D	0.72	0.36	0.24	0.18	0.14	0.12	0.1	0.09	0.08
	E	0.9	0.45	0.3	0.22	0.18	0.15	0.13	0.11	0.1
	F	1.08	0.54	0.36	0.27	0.22	0.18	0.15	0.13	0.12

Table 5. Insurance contracts in the market.

Limit	μ_L	σ_L	Frequency (π_L)	Expected Loss ($\pi_L.\mu_L$)	Premium (Exp. loss/limit) (%)
$500k	$200k	$125k	0.1	$20k	4
$1mn	$400k	$350k	0.15	$60k	6
$2mn	$1mn	$1mn	0.16	$160k	8
$5mn	$2.5mn	$1.25mn	0.2	$500k	10
$10mn	$4mn	$4mn	0.3	$1.2mn	12

Table 6. Insurance policies written by insurance panel.

	Policy count grouped by policy limit						
Insurer	$500k	$1mn	$2mn	$5mn	$10mn	Total exposure (mn)	Technical premium (mn)
Alpha	200	0	0	0	0	$100	$4.0
Beta	100	50	0	0	0	$100	$5.0
Charlie	50	20	15	5	0	$100	$7.1
Delta	30	0	5	5	5	$100	$9.9
Echo	0	0	0	0	10	$100	$12.0
Total	380	70	20	10	15	$500	$38.0

(1) Set severity distribution for each contract as in Table 5.
(2) Set frequency distribution as per Table 5—Poisson $\sim \pi_L.k$ where k is the number of each contract contained in the portfolio.
(3) Randomly sample the frequency of expected losses for each contract in the portfolio, to generate a number of losses for *each contract*, N_{loss}.
(4) Randomly sample from the severity distribution for each contract N_{loss} times, sum and record the losses.
(5) Run the above process 100 000 times.

The results of the simulations are presented in Table 8 (histograms of the generated loss distributions are provided in Fig. A1). The table contains the premium income for each insurer as previously determined, a capital level assumed to be held by the insurer equal to the average baseline loss in the simulation and reserves defined,

$$\text{Reserves} = \text{Premium Written} + \text{Capital}. \tag{25}$$

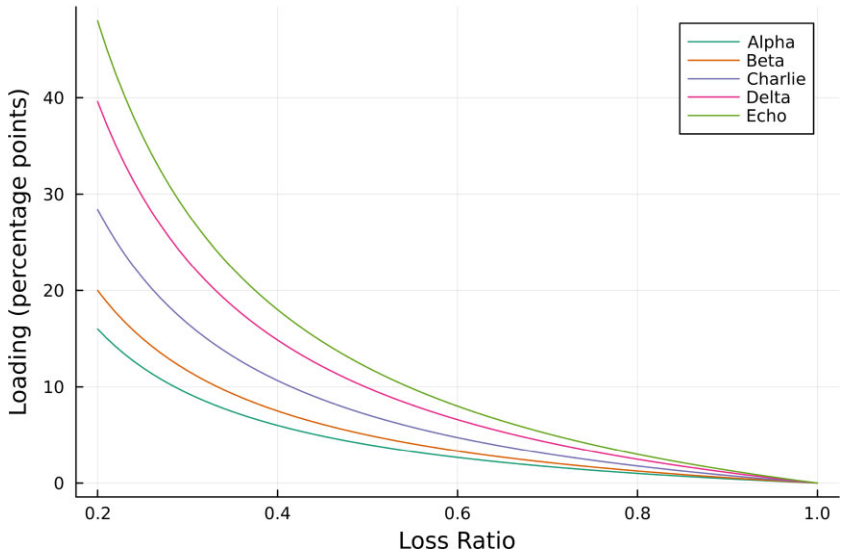

Figure 5. Loading versus target loss ratio for different insurance portfolios.

Table 7. Calculating premium loading rates for insurance companies based on simulated losses. The loading rate is expressed in percentage points.

Insurer	Technical Premium (mn)	Target loss Ratio (%)	Exposure (mn)	Technical Premium rate (%)	Loading (pp)	Weighted Average Charged Premium rate (%)
Alpha	\$4.0	50	\$100	4.0	4.0	8.0
Beta	\$5.0	50	\$100	5.0	5.0	10.0
Charlie	\$7.1	50	\$100	7.1	7.1	14.2
Delta	\$9.9	50	\$100	9.9	9.9	19.8
Echo	\$12.0	50	\$100	12.0	12.0	24.0

Table 8. Simulated losses.

	Assets			Losses			
Insurer	Premium Income (mn)	Capital (mn)	Reserves (mn)	Simulation Baseline average (mn)	Simulation Baseline SD (mn)	95% Stress test (mn)	97.5% Stress test (mn)
Alpha	\$8.0	\$3.6	\$11.6	\$3.6	\$0.8	\$8.2	\$9.4
Beta	\$10.0	\$4.4	\$14.4	\$4.4	\$1.3	\$13.6	\$17.4
Charlie	\$14.2	\$6.4	\$20.6	\$6.4	\$3.0	\$28.0	\$36.6
Delta	\$19.8	\$8.9	\$28.7	\$8.9	\$6.2	\$51.2	\$64.9
Echo	\$24.0	\$10.8	\$34.8	\$10.8	\$7.9	\$53.1	\$77.0

Along with the simulated loss values, we also calculate loss values for a 'stress test' type scenario, calculating the maximum loss in 95% and 97.5% of cases[15]. This is done via using the *quantile* function of *Distributions.jl* to calculate the respective frequency and severity at $F^{-1}(0.95)$ and $F^{-1}(0.975)$. The required values are then readily obtained. With these values obtained, we may now proceed to consider the interaction between reinsurance and the insurance portfolios.

15 These loss probabilities were selected rather than the Solvency 2 limit of 99.5% as it is assumed that an insurance company has loss tolerance of less than the formal Solvency 2 regulatory buffer. Insurers typically set internal risk appetite, via a process documented in, e.g. Lloyd's standard MR5 [71].

Considering the effect of capital

Suppose, as per Table 8 that the insurer has a capital buffer, which initially, is equal to the simulated average losses on its portfolio. We now examine the optimal reinsurance fraction which means the insurer would remain solvent in the event of losses of a specified magnitude. We consider ρ values for both the 95% and 97.5% stress tests. This means calculating the value of ρ which would set $\text{\$Profit}(L) = -K$ (Equation 23). The required expression is

$$\tilde{\rho}(L_{\text{stress}}) = \frac{(L_{\text{stress}} - \text{\$Premium Written} - K)}{L_{\text{stress}} - \text{\$Premium Written} + (\text{\$Exposure} \times \%\text{CC})}, \quad (26)$$

where %CC is the percentage ceding commission.

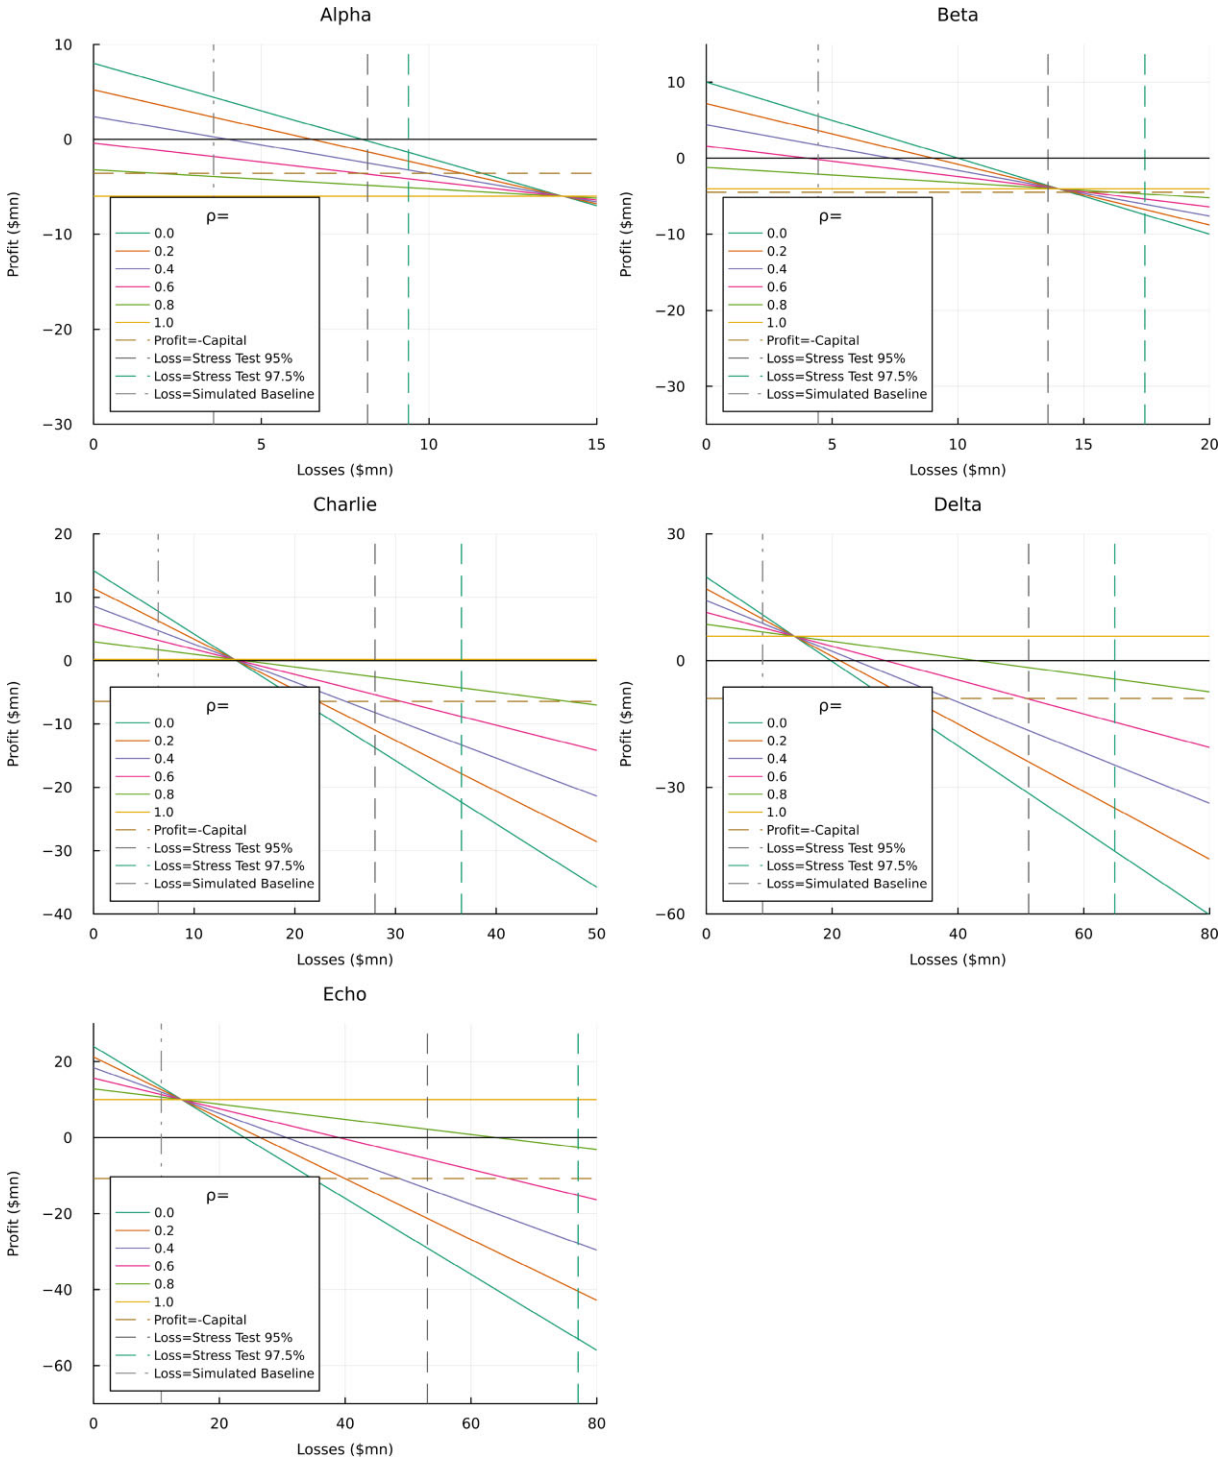

Figure 6. Insurer profit with varying quota share proportions (ρ). Note that the profit (y) axes in the charts are scaled differently for each subplot to allow the key features of the plots to be easily identified.

The solvency threshold for the insurer is Reserves $= L_{stress}$. If Reserves $> L_{stress}$ then we set $\bar{\rho} = 0$ as the insurer does not need reinsurance at this stress test loss level as it would remain solvent without it. For insurer Alpha, reserves exceed the stress test losses at both thresholds, while for Beta, reserves exceed only the 95% stress test loss. Figure 6 and Table 9 show the complete results of the analysis. Starting with Table 9, it appears that neither Alpha nor Beta should buy reinsurance. In the 97.5% stress-test, Beta is insolvent even with reinsurance. This suggests that Beta would need to implement a higher loading than that initially calculated to pass the stress test. For Charlie, Delta, and Echo, there is benefit in purchasing reinsurance as a quota share policy, though the optimal fractions appear fairly high. Consequently, the insurers might decide to buy less than the optimum but set capital higher. However, this then means that the market is not efficient. Table 9 also shows the profit each insurer would receive if *ex post* losses equal the simulated baseline with no

Table 9. Reinsurance ceding fractions that maintain insurer solvency at different stress-test values.

Insurer	Ceding Commission (%)	$\tilde{\rho}^{0.95}$	$\tilde{\rho}^{0.975}$	Profit if losses = simulation baseline ($mn)		
				$\rho = 0$	$\rho = \tilde{\rho}^{0.95}$	$\rho = \tilde{\rho}^{0.975}$
Alpha	−6.0	0.00	0.00	4.4	4.4	4.4
Beta	−4.0	0.00	0.87	5.6	5.6	−2.7
Charlie	0.2	0.53	0.71	7.8	3.8	2.4
Delta	5.8	0.60	0.71	10.9	7.8	7.3
Echo	10.0	0.47	0.67	13.3	11.7	11.1

reinsurance; with reinsurance at the $\rho^{0.95}$ fraction; and with reinsurance at the $\rho^{0.975}$ fraction. For Charlie, given that its weighted average premium rate is close to the reinsurer objective, it receives a scant ceding commission. Consequently, there is an opportunity cost of $4.0–5.2mn of purchasing quota share at the optimum relative to baseline simulated profit of $7.8mn. In a market where information is shared, there should not be an opportunity cost. For Delta and Echo, the purchase of quota share appears more attractive because of the more generous ceding commission. These are deliberately extreme examples, but in practice suggest that bargaining may occur between different insurers and reinsurers over the ceding commission, which introduces inefficiency into the market.

Figure 6 presents a more detailed picture of the simulations that yield the optimal ρ. For each insurer portfolio, we plot the insurer profit (Equation 23) as a function of losses for values of ρ between 0 and 1. The capital held (i.e. the average simulated loss as already discussed) is represented as a horizontal line and the average simulated losses for the 95% and 97.5% stress tests are represented as vertical lines. The intersection of the average simulated loss and the stress test allows for the optimal ρ to be read from the graphs. In the case of Alpha, it can be seen that on the $\rho = 0$ profit line, at the two stress test loss values (Points A and B), the profit exceeds the capital held. For reinsurance to be worth purchasing, the $\rho = 0$ profit line must be less than the capital horizontal lines at the stress test losses. Taking Echo as an example, with stress test loss of 95%, we can see that the horizontal capital and vertical loss lines intersect between the profit lines for $\rho = 0.4$ and $\rho = 0.6$ (Point C) As may be verified from Table 9, the reinsurance fraction for this case is 0.47. The comparable intersection for the 97.5% stress test (Point D) is at $\rho = 0.67$.

Excess of loss

Having considered the quota share case, it is worth considering the case of excess-of-loss insurance as an alternative to quota share for the insurers. Rather than using the capital buffer approach, we consider a simpler objective: that the insurer rather than holding a capital buffer buys insurance from a reinsurer to cover losses in excess of its cash premium income up to the limit of the 97.5% stress test loss value. To calculate the required parameters, we can use the simulated baseline losses already calculated in Simulation procedure. From these, we compute the number of instances of losses in the vector of generated losses that exceed the cash premium income but are less than the 97.5% stress test loss value.

The results are contained in Table 10. The portfolios of Alpha and Beta generate expected losses well below the level of cash premium income (see Appendix 1) and accordingly there is little benefit in excess-of-loss insurance. For Charlie, Delta, and Echo, it is interesting to note that the combined technical premium is $7mn. Recall that in Simulation 2: reinsurance supply and price, Distribution A in Table 4 gives the loss ratios for the reinsurer versus quoted premium for an expected $10mn of losses. If we assume that the reinsurer requires a loss ratio of 0.5 or better, then the minimum premium it will charge is 4%. For the insurance buyers, only for Echo is buying excess-of-loss reinsurance cheaper than buying quota share. Consequently, for each insurance portfolio, there is a different optimal reinsurance contract from the perspective of the insurance company seeking reinsurance.

The excess-of-loss premia calculated in Table 10 are computed using the individual joint distributions of frequency and severity for each of the five insurance companies. These are known only to each of those insurance companies alone and are not visible to the reinsurer. Consequently, there are information asymmetries between the insurers seeking reinsurance and the reinsurer. Insurers Delta and Echo know that the fair insurance premium rate for the excess-of-loss contracts specified in Table 10 are 5.9% and 6.6%, respectively. However, the reinsurer would offer these contracts at 4% premium rate based on its own distribution. Consequently, the insurers can, under these assumptions, buy reinsurance cheaper than its fair cost based on their avantageous knowledge of the 'true' distribution rather than the reinsurer's distribution which assumes simple log-normal distribution of a set of risks at a particular expected loss value. This illustrates how inefficiency and therefore financial imbalances between insurance and reinsurance may emerge as a consequence of different expected loss distributions, unlike in Table 2 where the reinsurer and insurer(s) had the same distribution of expected losses.

Simulation 3: insurance buyers of variable risk

We now consider a simulation in which buyers have heterogeneous preferences and risk tolerance. The interactions of the real insurance market are hard to model as insurance customers interact with insurance companies via insurance brokers who act as an intermediary. The flow of business is directed therefore partly by relationships (and so is not efficient in a traditional economic sense). However, it is possible to construct some simulations of insurance demand based on different characteristics and illustrate the utility demand model and how this may affect reinsurance pricing.

The insurance buyer faces a single utility maximization decision: for a given premium rate, how much cover does the agent with to purchase. This could be formalized in terms of expected utility (Equation 4) via variation of the risk aversion parameter, α, but this is not necessary for the example presented here. The insurance company must choose premium rates that it believes will not excessively deplete its capital for a certain level of risks, or plan to cede premium to reinsurance to cover that risk as demonstrated in the previous section. We will retain the contract limit structure from Table 5 for this analysis, meaning that insurance buyers choose one of the five contracts.

We will now assume that the more coverage the buyer takes, the more sophisticated its assessment of the risks are. This places a

Table 10. Excess of loss pricing example.

Insurer	XL reinsurance Coverage	Probability of Loss > cash premium (%)	Technical XL Reinsurance premium (mn)	QS Reinsurance Premium at $\hat{\rho}^{0.975}$ (mn)
Alpha	\$1.4mn xs \$8.0mn	0.0	\$0.0	\$0.0
Beta	\$7.4mn xs \$10.0mn	0.0	\$0.0	\$8.3
Charlie	\$22.4mn xs \$14.2mn	1.6	\$0.4	\$5.4
Delta	\$45.1mn xs \$19.8mn	5.9	\$2.7	\$3.6
Echo	\$53.0mn xs \$24.0mn	6.6	\$3.5	\$2.2

Table 11. Insurance buyer premium ceilings.

	Highest premium rate at which a buyer takes full coverage			Maximum number of customers		
Limit	Low risk (%)	Medium risk (%)	High risk (%)	Low risk	Medium risk	High risk
\$500k	14	20	26	46	46	46
\$1mn	13	18	23	32	32	32
\$2mn	12	16	20	16	16	16
\$5mn	11	14	17	8	8	8
\$10mn	10	12	14	4	4	4

constraint on the amount of loading the insurer can apply to the higher limit contracts. We will, as previously, fix the total *potential* cover available in the market at \$500mn and consider how this may be allocated among buyers. However, as will be illustrated, the risks associated with some contracts make them commercially unviable even if theoretically priceable. Table 11 sets out some arbitrary premia based on the subjective beliefs of the respective buyers, and the maximum number of contracts available in the market based on the overall capacity of \$500mn. We wish to stress that these numbers are established purely for model convenience and to illustrate the further difficulties to establishing efficiency under heterogeneous buyer beliefs. The assumption of market size is required to price potential reinsurance on insurer policies.

For this analysis, we set the expected severity loss mean equal to a quarter of the policy limit and the standard deviation to half the mean. Unlike in the previous section, we will allow the distribution of expected losses to vary with different clients and have a mixture of buyers considered low-, medium-, and high-risk with different distributions accordingly. We assume that the variation in risk characteristics of the three buyer groups is expressed through variation in frequency.

We assume that reinsurers consider the risks involved for the three different risk categories and apply distributions A, C, and E (Table 3) to low, medium and high risks effectively, and target loss ratios of 0.3, 0.5, and 0.7, respectively. This means that the reinsurance charges for the portfolios are 6%, 11%, and 13%.

We now consider the distributions associated with the different contracts. Table 12 shows the severity and frequency distributions for each policy. We have fixed the severity on each contract and assumed that riskier clients have a higher expected frequency of claims. This assumption could, of course, be varied further, but this approach suffices for the purposes of this example. From this, we simulate the losses with 100 000 runs and derive the expected loss for the entire set of possible contracts. This is shown in Table 13 along with the expected average loss per contract derived using the assumptions for maximum number of customers outlined in Table 11.

With this calculated, we can then derive the technical premium for each contract, which is shown in Table 14. Comparing with Table 11, we can see that for the \$5mn and \$10mn limits, the high risk technical premium is higher than what customers are willing to pay. It may be possible in this case for the insurer to instigate a deductible and reduce the premium. Otherwise, margin is very limited for medium-risk \$5mn and \$10mn limits, which might also motivate introducing a deductible.

We now consider the capital requirements associated with the insurance policies. Table 15 shows the expected losses for $F^{-1}(0.995)$ and $F^{-1}(0.5)$ for frequency and severity respectively for both the whole set of contracts and also per contract. Each insurer must decide how to allocate its available capital and how much reinsurance to purchase. Rather than calculating sample portfolios, we will simply calculate the reinsurance fraction that is optimal based on Equation (23).

Based on the stress test loss values, and assuming that the insurer writing each contract holds capital equal to the expected value of losses for the contract (Table 13), we can then derive the optimal reinsurance fraction for each contract (Table 16). As in the prior section (Table 9), this is calculated by calculating the reinsurance fraction that sets the profit to the insurer equal to $-K$, i.e. at the level of loss given in the stress test, the insurer breaks even if it holds this proportion of reinsurance. As the buyers of the smaller contracts are less knowledgeable and will accept a higher premium, the reinsurance fraction is lower as the insurer writes more premium. However, the reinsurance fraction increases from an average of 20% for the \$500k limit contract to as high as 64% for the medium-risk \$10mn limit contract. It is clear from this analysis that while it is possible to achieve risk transfer between insurance buyer, insurance company and reinsurer, for a simulated market, achieving convergence of distributions is extremely unlikely as each party is incentivized to maximize their profit rather than target efficiency.

We have stopped short of simulating the allocation of policies to individual insurers as to model competitive market dynamics under uncertainty with heterogeneous beliefs is a complex problem that in itself might fill multiple papers. However, it is hoped that the simulation presented illustrates the additional dynamics that heterogeneous buyer beliefs brings to the challenges of modelling cyber-insurance and reinsurance. To place the simulation results in context with the

Table 12. Distribution specification for insurance contracts offered to buyers.

	Severity			Frequency, Poisson(λ)		
Limit	μ_L	σ_L	Distribution	Low risk	Medium risk	High risk
$500k	$125k	$62.5k	LogNormal(11.6,0.22)	4.6	11.5	23
$1mn	$250k	$125k	LogNormal(12.3,0.22)	6.4	12.8	19.2
$2mn	$500k	$250k	LogNormal(13.0,0.22)	4	8	12
$5mn	$1.25mn	$625k	LogNormal(13.9,0.22	2	4	6
$10mn	$2.5mn	$1.25mn	LogNormal(14.6,0.22)	1	2	3

Table 13. Expected losses for policies.

	Expected Loss (Total, $mn)			Expected loss per contract ($k)		
Limit	Low risk	Medium risk	High risk	Low risk	Medium risk	High risk
$500k	0.53	1.32	2.63	11	29	57
$1mn	1.47	2.93	4.40	46	92	138
$2mn	1.84	3.67	5.50	115	229	344
$5mn	2.28	4.59	6.89	285	574	861
$10mn	2.30	4.59	6.86	574	1,147	1,716

Table 14. Technical premium for insurance contracts.

	Technical premium (%)		
Limit	Low risk	Medium risk	High risk
$500k	2.3	5.7	11.5
$1mn	4.6	9.2	13.8
$2mn	5.7	11.5	17.2
$5mn	5.7	11.5	17.2
$10mn	5.7	11.5	17.2

US cyber-insurance market, in 2020, according to the NAIC [72], there were approximately 4 million cyber-insurance policies written in the US market, with the top 20 insurers taking 68% market share. The report for 2021 does not provide a policy number, but notes that almost 50% of cyber-insurance premia were ceded.

Discussion

The simulations show the difficulty of achieving economic efficiency in an artificial cyber-insurance market even using relatively standard distributions and contract structures. However, as has been stressed, just because a market is not efficient does not mean that transactions cannot take place. We now consider some of the further informational barriers to facilitating smooth transfer of cyber-risk. Issues of data transparency, incident measurement, and reporting—making relevant data publicly available—are particularly crucial in enabling agents to make informed pricing decisions.

Information asymmetry

By and large insurance and reinsurance companies operate in environments where high quality precision signals about loss risks exist. For example, in the case of natural catastrophes, their frequencies are well known and established over many periods. Further, there are enough tail events to help construct reasonable approximations of extremes. When it comes to events regarding human interactions, such as crime, illness, death, or accidents, these are reported by statute to the relevant central authorities. This data is publicly available. In both these cases agents at all levels share the public signals and can condition their private expectations on good quality evidence. Of course, there may be variability in the accuracy of private expectations based on individual interpretation of the data or circumstances. This set-up allows the buyers of insurance the calculate their expected loss in a well informed manner and the insurance companies, based on the public information, can quote a premium. In turn the reinsurers share the same beliefs as no further information is available to them regarding the likelihood of the different states of nature.

When it comes to cyber-risk and cyber-insurance, the state of data curation and sharing is far more nascent than for other insurance perils and it is reasonable to argue that there is no high quality public signal to inform all agents' priors. In the regulation of the aviation industry, it is standard to require reporting of 'near misses' so that lessons can be learnt and procedures updated to lessen the risk of future accidents. It is possible that this might be addressed by vendor telemetry—an insurer might have a series of recommended cyber-security solution providers that their clients could sign up for as part of their insurance package who would share data with the insurer. This raises potential issues of confidentiality.

Cyber-insurer loss experience

The United States National Association of Insurance Commissioners publishes an annual report on the cyber-insurance market derived from its Property/Casuality Annual Statement [72]. Table 17 presents this information for the four years currently available. In 2018 and 2019, the data was presented separately for standalone and package policies but in 2020 and 2021 was presented for combined policies. We have adjusted for this to present the data on a consistent basis.

It is notable that the ransomware epidemic from 2020 to 2021 had a marked effect on experienced loss ratios for some insurers[16]. However, there are pockets of differentiation. For example, the Hartford Insurance Company specializes in insurance for smaller companies, creating a fairly well diversified portfolio of insurance contracts where the holders are unlikely to fall victim to sophisticated, targeted

16 This has been widely reported in the trade press—see, e.g. [73].

Table 15. Stress test losses for policies, with frequency set at $F^{-1}(0.995)$, severity at $F^{-1}(0.5)$.

	Stress test loss (total, $mn)			Stress test loss per contract ($k)		
Limit	Low risk	Medium risk	High risk	Low risk	Medium risk	High risk
$500k	1.2	2.3	4.0	27	51	87
$1mn	3.1	5.1	6.9	98	161	217
$2mn	4.5	7.2	9.8	280	447	615
$5mn	6.7	11.2	14.5	839	1398	1817
$10mn	8.9	13.4	17.9	2236	3354	4472

Table 16. Optimal reinsurance purchase fraction for each contract implied by stress test values.

	Optimal reinsurance fraction for each contract		
Limit	Low risk	Medium risk	High risk
$500k	0.23	0.23	0.20
$1mn	0.31	0.30	0.25
$2mn	0.41	0.40	0.36
$5mn	0.51	0.53	0.48
$10mn	0.63	0.64	0.60

ransomware attacks given the potential revenue available. For these companies, basic defences and security software should help mitigate against losses. Figure 7 plots the losses experienced in the underwriting year versus the premium written and a linear trend line with intercept fixed at 0. The slope of the fitted trend line is then the loss ratio. The average loss ratio remained fairly stable across the two years, but it is striking that less than 30% of premia received was, on average, retained by the underwriting insurer. The aforementioned NAIC report states that some 50% of premia for cyber-insurance was ceded to the reinsurance market. There is some evidence to support the premise of a disconnect between expected and experienced losses in cyber-insurance pricing. Woods *et al.* (2021) [62] develop a distribution of cyber-losses based on insurance company filings in the USA. They note that their model significantly underpredicts losses in relation to *ex post* losses reported in other literature. The underpricing of premia implies that either

- Insurers believe they can diversify loss risk.
- Customers were not willing to pay the technical premium and insurers are pursuing a 'loss-leader' strategy.

The entry of Arch Insurance also merits comment. Arch insurance provides capacity[17] to a relatively new managing general agent, Coalition Inc., providing 'active cyber-insurance'. Active cyber-insurance is a relatively new product, which merges the roles of an outsourced security provider and a traditional cyber-insurer. This reduces some of the risks of asymmetric information transfer associated with cyber-insurance from the perspective of the insurer. The trade-off between cyber-insurance and security investment has been modelled by Mazzoccoli and Naldi (2020) [74] and Skeoch (2022) [75].

Comparison to simulated results

Comparing the experienced losses by insurance companies, our assumption regarding the adoption by reinsurance firms of their own private distributions for both severity and frequency of successful cyber-incidents and subsequent losses at this stage of development of this nascent market seems well-grounded on the available evidence.

The evolution of proportional losses across 15 major insurance companies over the period 2018–2021 presented in Table 17 reveals a somewhat unstable path. Both the average loss and its distribution exhibits both wide variability and an increasing trend. Specifically in 2018, average losses were 25.3% of the premia collected and this measure has monotonically increased to 68.3% by 2021. At the same time the maximum losses have more than doubled from 57% to 130% by 2021. The cross-sectional standard deviations exhibit the same monotonic trended pattern.

Attempting to fit a log-normal distribution over the whole period for the companies in the sample using the same methodology for fitting such distributions in the simulations shows that the kernel[18] of the empirical distribution deviates significantly from the normal and reveals slight bimodality (Fig. 8). It is also notable that the fitted distributions underestimates the tail of large losses, which is arguably a significant consideration for reinsurance companies.

Faced with such movements of the cross sectional distributions, meaningful aggregation of the losses experienced by individual insurance companies does not seem effective. In the light of this (admittedly cursory) review of the statistical evidence presented in this paper, Assumption 4 in Introducing reinsurance seems justified.

Loss transparency

We consider what happens if agents only selectively claim on losses from an insurer. In an insurance analysis, it is usually assumed that every agent is aware of the incidents they experience. This is a reasonable assumption for some categories of cyber-incidents, such as ransomware, although other cyber-incidents such as data breaches might not be detected until some time after the event. Agents report some incidents to an insurer and thus a claim is made; some incidents go undisclosed (in insurance, this is known as IBNR—incurred but not reported). More formally, at time t, the agent may be aware of the indicent and its damage so the state of the world in which the incident occurs, s is known to them. The agent might inform the insurer about the state so the insurance knowledge of the state s is conditional on the revelation of the agent. Now, the insurer knows that their distribution is not the objective one but only a partial revelation due to the agents selectively choosing to report losses. The insurer then tries to approximate the objective distribution but it will be with error. In the event that reinsurers know that different insurers have different approximations of the true distribution, they will use some kind of averaging across these approximations to quote reinsurance premiums. The results are:

- No insurer is offered a fair premium given their approximation of the true distribution.

17 https://www.coalitioninc.com/en-ca/announcements/Arch-Insurance-Backs-Coalition-With-Long-term-Capacity-Across-Cyber-Insurance-Programs

18 See Epanechnikov (1969) [76].

Table 17. Cyber-insurer loss experience in the US market (†denotes weighted average by DWP).

	Direct written premium ($mn)				Loss ratio			
Firm	2018	2019	2020	2021	2018† (%)	2019† (%)	2020 (%)	2021 (%)
CHUBB LTD GRP	320.73	355.28	404.14	473.07	28.6	27.7	61.0	76.9
FAIRFAX FIN GRP	38.15	65.01	108.69	436.45	23.4	51.6	55.7	51.9
AXA INS GRP	255.87	229.68	293.03	421.01	57.2	65.7	98.2	86.5
TOKIO MARINE HOLDINGS	44.59	46.91	78.16	249.79	30.6	17.1	51.1	43.8
AMERICAN INTL GRP	232.31	226.20	228.42	240.61	36.1	55.4	100.6	130.6
TRAVELERS GRP	146.23	178.53	206.82	232.28	22.4	32.1	85.5	72.7
BEAZLEY INS CO INC	110.95	150.94	177.75	200.88	7.8	22.0	47.9	38.7
CNA INS GRP	83.36	94.72	119.61	181.38	26.9	33.2	105.7	87.5
ARCH INS GRP	–	–	–	171.94	–	–	–	9.2
AXIS CAPITAL GRP	76.00	97.31	133.55	159.06	7.2	18.5	46.2	105.2
ZURICH INS GRP	43.32	43.67	64.43	151.87	18.2	86.9	40.4	76.9
LIBERTY MUT GRP	66.50	68.38	41.86	138.22	38.9	23.3	30.0	95.2
SOMPO GRP	34.05	49.71	72.59	133.52	56.7	29.3	114.1	54.3
BCS INS GRP	69.50	76.06	86.58	132.04	10.4	32.9	59.1	80.1
HARTFORD FIRE 7 CAS GRP	39.70	49.74	102.86	123.16	16.4	31.6	25.4	16.3

Source: NAIC, Researcher calculations.

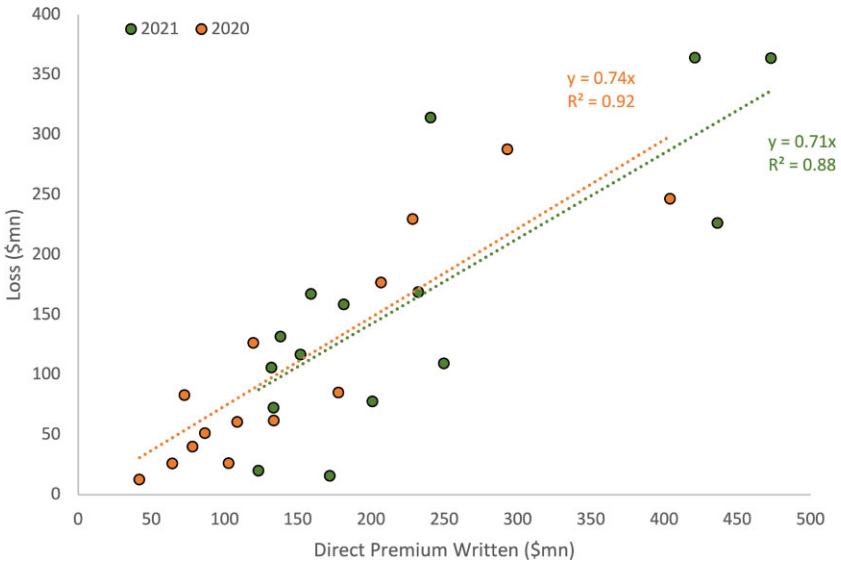

Figure 7. US cyber-insurer losses vs premium written.

- No agent is offered a fair premium as the insurance offer is based on a distribution different to their own.
- Objectively measured data is absent at all levels because reporting is a choice.

Consistency of reference

There is a significant problem with the standard actuarial modelling cycle approach to cyber-insurance: the evolution of systems over time, which is quite unique in its complexity in relation to other perils. Calibration of models using events such as WannaCry have poor future predictive power as the security vulnerabilities it exploited have been patched, Windows XP is less widespread than it was and the operating systems that replaced it have better, though of course not perfect, security by design. In economics, this can be couched in clients' Bayesian updating of their distributions; they do not and cannot observe incidents of other clients (other than indirectly via media reports) so there is no need to converge to a stationary distribution at the client level. The consequence of this is that the insurers and reinsurers may have a better understanding of the fair price of risk, but buyers do not share the same concern and thus are not willing to pay the demanded premium for the insurance.

Supply and demand

In the insurance industry, it is common to describe the state of the market as 'hard' or 'soft'. In a soft market, supply exceeds demand placing downward pressure on premium, whereas in a hard market the converse is true. Often the experience of losses in a particular class of business will result in a market hardening. This has important implications for the pricing of cyber-insurance by a vendor. In a soft market, the insurer must charge the lowest premium it can actuarially justify to build market share. In a hard market, the insurer should charge the highest realistic premium possible. If the market were efficient, it would converge to some form of equilibrium but if not it may swing between financial imbalances. There is evidence that in the early stages of the cyber-insurance industry, some insurers operated a very experimental approach to pricing. Woods (2023) [77]

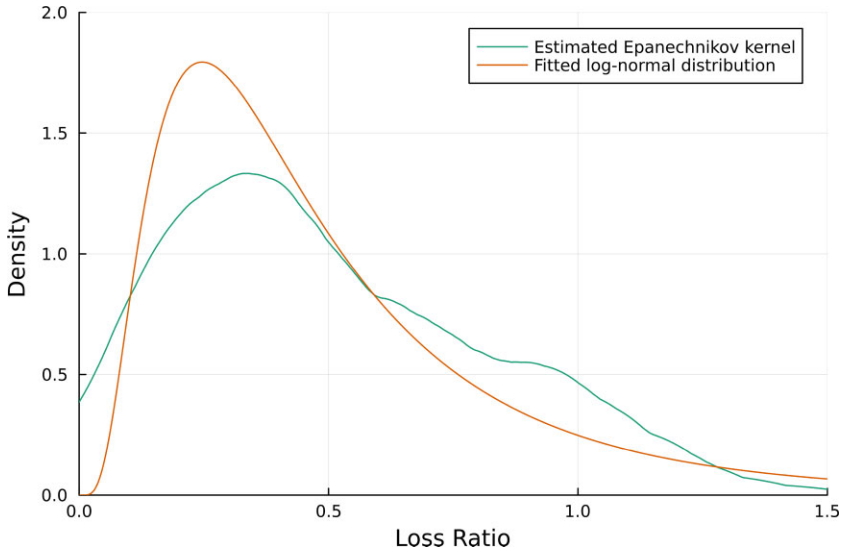

Figure 8. Epanechnikov kernel versus fitted log-normal distribution for NAIC reported cyber-insurance loss ratios, 2018–2021.

provides an account of one large US insurer, AIG, whose Chief Operating Officer admitted that their early cyber-insurance models were a 'complete guess'. The same insurer then suffered loss ratios of 100% and 130% in 2020 and 2021, respectively (Table 17), suggesting that even if refined and updated, the pricing models may have underestimated the claim frequency or severity.

Further work

We have considered simulations in which losses are uncorrelated. An interesting next step would be to consider the correlation of losses and implement the modeling strategy presented in this paper using more complex loss-generating functions, such as those reviewed in Actuarial models, than the simple joint distributions of severity and frequency used in this paper. It would also be instructive to compare the results of simulations of distributions proposed by Eling *et al.* (2019) [32] and Woods *et al.* (2021) [62], with insurer loss data. Claims data is deeply confidential to insurance companies, however, so the results of such analysis would unlikely be able to be widely disseminated unless extensively anonymized.

In the simulations, we focused on the supply dynamics of insurance and in particular the interaction between insurers and reinsurers. The model provides for consideration of buyer preferences, which at this stage we have explored only briefly in the first simulation to illustrate how buyer utility can affect coverage. A further piece of work would be to explore the price sensitivity of buyers of insurance coverage and how these preferences propagate through the information chain to reinsurers. The model simulations considered only a small number of market participants; with a more complex set of interactions, it may be possible to attempt to determine the optimal market size by introducing appropriate constraints and incentives. A further addition might be to consider multiple reinsurers with different risk tolerances; however, this would add considerable complexity to the model and is outside the scope of the framework introduced in this paper.

Conclusions

This paper has developed an artificial yet realistically structured model of the cyber-insurance market considering all three levels of agent interactions. The model incorporates the demand choices of the consumers/buyers of cyber-insurance, their suppliers—insurance companies offering contracts—and reinsurance companies providing additional underwriting capacity.

The extent to which an insurance market facilitates smooth risk transfer is linked to the sharing of information by participants regarding the distribution of losses. We argue that this condition is very unlikely to hold in the cyber-insurance market. Disagreements on loss expectations means that cyber-insurance contact pricing will be considered inefficient at both the retail and wholesale levels, leading to lower societal benefit. The purpose of this paper was to quantify such inefficiency within the confines of a three-tier market under miscellaneous types of disagreements in loss expectations among the participants at each tier.

To establish a benchmark to gauge the extent of inefficiency, we have simulated a simple market where all agents share a distribution of losses based on two loss frequencies. From this simulation, we obtained the 'efficient' measures of reinsurance premium and the proportional participation of reinsurers. We found that simulated loss reduction to the insurers is almost identical to the cost of reinsurance (bar small statistical errors), as expected. This case represents the economically efficient market outcome.

Maintaining all the behavioural parameters from the first simulation, we then proceeded to compute expected losses and reinsurance premiums based on diverse distributions held by insurance companies and reinsurers. Both insurers and reinsurers independently price premiums to meet target loss ratios based on distinct and subjective distributions. Under conditions where losses are close to the modal simulated value, insurers are typically not incentivized to buy reinsurance. However, when considering relatively extreme losses under a 'stress test' type scenario, the value of reinsurance emerges to some insurers whose distributions are relatively heavy tailed in comparison to others. For such insurers, the upfront cost of such reinsurance is justified by the avoidance of ruin under high loss scenarios.

Even within the confines of this simple example, the divergence in distributions, expectations, and objectives demonstrates that efficient pricing is hard to achieve. It should be noted that whilst there are specialists in cyber-insurance operating within the reinsurance market, cyber-insurance itself competes with other lines of insurance for allocation of specialty reinsurance capital. Based on this, we used a uniform cost of reinsurance in the second of our two simulations. This is the outcome of the reinsurer holding a private loss distribu-

tion. This condition may reduce the reinsurance capital allocated to cyber-insurance. It is notable that, according to industry sources [2], there are only a limited range of nonproportional reinsurance structures available to the cyber-insurance market; in other words, the vast majority of written reinsurance is quota share. This implies significant uncertainty among tail risks by reinsurers at the present time.

Our findings suggest that the cyber-insurance market will continue to face potential financial imbalances. That is, it will be highly profitable for some participants and costly for others. This is already evident in data on cyber-insurer loss data (Table 17). There has been considerable progress in the academic literature on theoretical modelling of cyber-losses and on empirical analysis. However, access to reliable and transparent data remains a problem for researchers as insurance claims data is confidential and highly guarded. Kasper (2019) [78] has developed a model for evaluating the feasibility of cyber-catastrophe bonds, while more recently Braun *et al.* (2023) [79] have noted that an insurance-linked securities market to support cyber-insurance may struggle to develop without better cyber-modelling. It is likely that this will be a high priority for the market, based on an extensive report by the Geneva Association (2023) [7]. There have been some efforts in the literature to move towards improving cyber-modelling specifically from an insurance perspective, such as Woods *et al.* (2021) [62] and Kasper and Grossklags (2022) [80]. As the literature develops, the model introduced in this paper may provide a convenient framework for evaluating more intricate distributions than the standard ones used for the simulations in this paper. Without a means of accessing reliable data on cyber-losses, insurance buyers will have to continue to form highly subjective probability distributions. In a recent paper, Bajoori *et al.* (2022) [81] argue for the creation of an official registry of cybersecurity experts with a duty to report, which has also been proposed by the UK Government[19]. Such public data might allow for the creation of distributions of cyber-losses and help contribute to reducing information asymmetries.

The cyber-insurance market is still at as stage of relative infancy. The current institutional setup does not appear fully conducive to the delivery of efficient market outcomes at this juncture. Achieving efficiency requires commonly held beliefs and stationary loss distributions. Whether such conditions can be achieved and maintained is questionable given the dynamic nature of cyber-threats. Our provisional conclusions are that the most likely market structure will involve firms specializing in particular insurance contracts covering different ranges of loss limits, with varying access to reinsurance based on these contracts. The overall outcome will be that the capital capacity of this market will be below its optimal size under shared informational conditions.

Acknowledgements

H.S. was employed by Convex Insurance between July and August 2021, which provided the inspiration for part of this work. He would like to thank Christophe Chandler, Rob Smart, and Harry Thompson in particular for helpful discussions during this time and for sharing their knowledge and acumen on the reinsurance markets. The arguments and analysis presented in this research are solely those of the authors and should not be interpreted as representative of the views or business practices of Convex Insurance, its clients or its counterparties. H.S and C.I. would like to extend particular thanks to Martin Eling for most helpful comments and recommendations that have greatly improved the structure and substance of the arguments within this paper. H.S. would like to thank Marie Vasek for helpful discussions as well as his fellow PhD candidates at UCL for their support and encouragement throughout this work. Anonymous reviewers for the Workshop on the Economics of Information Security 2023 provided helpful feedback and suggestions. The anonymous reviewers for the *Journal of Cybersecurity* provided extensive comments and suggestions that greatly improved the paper. This work was supported partly by the UK Engineering and Physical Sciences Research Council grant for Doctoral Training EP/R513143/1.

Author contributions

Henry R. K. Skeoch (Conceptualization, Methodology, Project administration, Software, Visualization, Writing – original draft), and Christos Ioannidis (Methodology, Supervision, Writing – original draft, Writing – review & editing).

Conflict of interest

This work was originally undertaken while Henry Skeoch was a PhD candidate at UCL. He is now employed at Beazley Group, who are an active participant in the cyber-insurance market, and while revisions to this paper have been undertaken since this employment commenced, the core conclusions and analysis contained within the paper were formed prior to his employment by Beazley. The views expressed in this paper are solely those of Henry Skeoch in a scholastic capacity and do not represent the views of Beazley. None of the content of this paper should be read as representative of Beazley's business practices, views on the market, or interactions with clients, brokers, and other counterparties.

19 https://www.ncsc.gov.uk/information/ncsc-assured-cyber-security-consultancy

References

1. Aon PLC. U.S. Cyber market update: 2022 U.S. cyber insurance profits and performance. 2022. https://www.aon.com/getmedia/438dfae5-3004-4f60-9698-d85fb6770868/20230920-2022-us-cyber-market-update.pdf (13 February 2024, date last accessed).
2. Guy Carpenter. Through the looking glass: interrogating the key numbers behind today's Cyber market. 2023. https://www.guycarp.com/content/dam/guycarp-rebrand/pdf/Insights/2023/Guy_Carpenter_Cyber_(Re)insurance_Market_Report_Publish_rev%20.pdf (13 February 2024, date last accessed).
3. Adam M Josefs M, Ashworth S. *et al.* Cyber risks in a new era: reinsurers could unlock the Cyber insurance market. 2021. https://www.spglobal.com/ratings/en/research/articles/210929-cyber-risks-in-a-new-era-reinsurers-could-unlock-the-cyber-insurance-market-12118547 (18 December 2023, date last accessed).
4. Gallagher Re. Cyber in the 2020s: a question of capacity. White Paper, 2021. https://www.ajg.com/gallagherre/-/media/files/gallagher/gallagherre/cyber-capacity-whitepaper.pdf (18 December 2023, date last accessed).
5. Gallagher Re. The future of cyber (Re)insurance. 2022. https://www.ajg.com/gallagherre/-/media/files/gallagher/gallagherre/future-of-cyber-reinsurance.pdf (18 December 2023, date last accessed).
6. Brew O., The all risk cyber challenge. Kansas City: Lockton Re, 2023. https://global.lockton.com/re/en/news-insights/lockton-re-cyber-report-says-market-needs-cyber-product-clarity (18 December 2023, date last accessed).
7. Pain D., Cyber risk accumulation: fully tackling the insurability challenge. Zürich: The Geneva Association, 2023. https://www.genevaassociation.org/sites/default/files/2023-11/cyber_accumulation_report_91123.pdf (18 December 2023, date last accessed).
8. Lloyd's of London. How the market works. 2022. https://www.lloyds.com/about-lloyds/our-market/lloyds-market (18 December 2023, date last accessed).
9. Institute and Faculty of Actuaries. https://actuaries.org.uk/membership/types-of-membership/ (18 December 2023, date last accessed).
10. Albrecher H, Beirlant J, Teugels JL., *Reinsurance: Actuarial and Statistical Aspects*. Hoboken: John Wiley & Sons, 2017.
11. Kiln R., *Reinsurance Underwriting*. Oxfordshire: Taylor & Francis, 2017.

12. Kurz M., *The Market Power of Technology: Understanding the Second Gilded Age*. New York: Columbia University Press, 2023.
13. Association of British Insurers. What does cyber insurance cover?. https://www.abi.org.uk/products-and-issues/choosing-the-right-insurance/cyber-insurance/what-does-cyber-insurance-cover/ (18 December 2023, date last accessed).
14. AIR worldwide. AIR Estimates Losses for the Marriott Breach Will Be Between USD 200 Million and USD 600 Million. 2018. https://www.air-worldwide.com/news-and-events/press-releases/AIR-Estimates-Losses-for-the-Marriott-Breach-Will-Be-Between-USD-200-Million-and-USD-600-Million/ (18 December 2023, date last accessed).
15. Reuters. Meta Platforms must face medical privacy class action. 2023. https://www.reuters.com/legal/meta-platforms-must-face-medical-privacy-class-action-2023-09-08/ (18 December 2023, date last accessed).
16. Devanny J, Martin C, Stevens T. On the strategic consequences of digital espionage. *J Cyber Pol* 2021;**6**:429–50.
17. Lloyd's Underwriting and Investment Phrases Glossary. https://www.lloyds.com/about-lloyds/investor-relations/financial-performance/financial-results/glossary (18 December 2023, date last accessed).
18. Business Insurance. Concerns over scope of cover limit cyber reinsurance capacity. https://www.businessinsurance.com/article/00010101/NEWS06/912360092/Concerns-over-scope-of-cover-limit-cyber-reinsurance-capacity (18 December 2023, date last accessed).
19. Cyber reinsurance must triple by 2030, capital markets are key: Howden. https://www.artemis.bm/news/cyber-reinsurance-must-triple-by-2030-capital-markets-are-key-howden (18 December 2023, date last accessed).
20. Munich Re prepared to give up cyber business over accumulation concerns. https://www.theinsurer.com/reinsurancemonth/golling-munich-re-prepared-to-give-up-cyber-business-over-accumulation-concerns/ (18 December 2023, date last accessed).
21. Nowak P, Romaniuk M. Pricing and simulations of catastrophe bonds. *Insur Math Econ* 2013;**52**:18–28.
22. Cummins JD. Cat bonds and other risk-linked securities: state of the market and recent developments. *Risk Manage Ins Rev* 2008;**11**:23–47.
23. Barrieu P, Albertini L. *The Handbook of Insurance-Linked Securities*. Hoboken: Wiley Online Library, 2009.
24. Dionne G, Harrington SE. *Foundations of Insurance Economics: Readings in Economics and Finance*. Vol. **14**. Berlin: Springer Science & Business Media, 2013.
25. Borch K. Equilibrium in a reinsurance market. *Econometrica* 1962;**30**:424–44.
26. Schlesinger H, Doherty NA. Incomplete markets for insurance: an overview. *J Risk Insur* 1985;**52**:402–23.
27. Froot KA, O'Connell PG. The pricing of US catastrophe reinsurance. In: *The Financing of Catastrophe Risk*. Chicago: University of Chicago Press, 1999, 195–232.
28. Bessy-Roland Y, Boumezoued A, Hillairet C. Multivariate Hawkes process for cyber insurance. *Ann Actuar Sci* 2021;**15**:14–39.
29. Hillairet C, Réveillac A, Rosenbaum M. An expansion formula for Hawkes processes and application to cyber-insurance derivatives. *Stoch Proces Appl* 2021;**160**:89–119.
30. Hillairet C, Lopez O. Propagation of cyber incidents in an insurance portfolio: counting processes combined with compartmental epidemiological models. *Scand Actuar J* 2021;**2021**:671–94.
31. Biener C, Eling M, Wirfs JH. Insurability of cyber risk: an empirical analysis. *Geneva Pap R I-Iss P* 2015;**40**:131–58.
32. Eling M, Wirfs J. What are the actual costs of cyber risk events?. *Eur J Oper Res* 2019;**272**:1109–19.
33. Baldwin A, Gheyas I, Ioannidis C. *et al.* Contagion in cyber security attacks. *J Oper Res Soc* 2017;**68**:780–91.
34. Bessy-Roland Y, Boumezoued A, Hillairet C. Multivariate Hawkes process for cyber insurance. *Ann Actuar Sci* 2021;**15**:14–39.
35. Arrow K. Aspects of the theory of risk-bearing. Helsinki: Yrjo Jahnssonin Saatio, 1965.
36. Debreu G., Theory of value: an axiomatic analysis of economic equilibrium, Vol. **17**. New Haven: Yale University Press, 1959.
37. Starr RM., Optimal production and allocation under uncertainty. *Q J Econ* 1973;**87**:81–95.
38. Feiger G., Diverse anticipations, rational anticipations, ex ante efficiency and ex post efficiency. Stanford: Graduate School of Business, Stanford University, 1976.
39. Simon SI. The dilemma of war and military exclusion clauses in insurance contracts. *Am Bus LJ* 1981;**19**: 31.
40. Woods DW, Weinkle J. Insurance definitions of cyber war. *Geneva Pap R I-Iss P* 2020;**45**: 639–56.
41. Rovetto Jr JM., Cyberwarfare & cyber insurance: exploring when a cyberattack can negate a cyber insurance claim. *J Bus Tech L* 2022;**18**: 309.
42. Brunner I., Insurance policies and the attribution of cyber operations under international law: a commentary. *NYUJ Int'l L Pol* 2022;**55**: 179.
43. Lyons Hardcastle J., Insurers can't use 'act of war' excuse to avoid Merck's $1.4B NotPetya payout. Princeton: BankInfoSecurity, 2023. https://www.theregister.com/2023/05/03/merck_14bn_insurance_payout_upheld/ (18 December 2023, date last accessed).
44. Chaudhry T., State backed cyber-attack exclusions. Market Bulletin, https://assets.lloyds.com/media/35926dc8-c885-497b-aed8-6d2f87c1415d/Y5381%20Market%20Bulletin%20-%20Cyber-attack%20exclusions.pdf (18 December 2023, date last accessed).
45. Cyber war clauses. 2023. https://www.lmalloyds.com/LMA/Underwriting/Non-Marine/Cyber_Clauses/cyber_war_clauses.aspx#:~:text=This%20clause%20includes%20a%20writeback,and%20manage%20the%20limits%20offered. (18 December 2023, date last accessed).
46. Wolff J. The role of insurers in shaping international cyber-security norms about cyber-war. *Contemp Secur Pol* 2023;**45**:1–30.
47. Kurz M. On rational belief equilibria. *Econ Theor* 1994;**4**:859–76.
48. Kuhn TS., The structure of scientific revolutions. In: SK Thomas (ed.), *International Encyclopedia of Unified Science*, Vol. **2**, 2nd edn. Chicago: University of Chicago Press; 1970.
49. Hammond PJ. Ex-ante and ex-post welfare optimality under uncertainty. *Economica* 1981;**48**:235–50.
50. Rees R, Wambach A. The microeconomics of insurance. *Found Trends Microecon* 2008;**4**:1–163.
51. Finetti De B. Theory of probability. a critical introductory treatment. In: *Wiley Series in Probability and Mathematical Statistics*. London: Wiley, 1974.
52. Skiadas C., Smooth ambiguity aversion toward small risks and continuous-time recursive utility. *J. Polit Econ* 2013;**121**:775–92.
53. Miccolis RS. On the theory of increased limits and excess of loss pricing. *PCAS LXIV* 1977;**27**:8085483.
54. Williams D. *Probability With Martingales*. Cambridge: Cambridge University Press, 1991.
55. Moore GE., Moore's Law at 40, *Understanding Moore's Law*. Chemical Heritage Press, Philadelphia, PA, 2006.
56. Benson CL, Magee CL., Quantitative determination of technological improvement from patent data. *PloS ONE* 2015;**10**:e0121635.
57. Funk JL, Magee CL., Rapid improvements with no commercial production: how do the improvements occur?. *Res Pol* 2015;**44**:777–88.
58. Nagy B, Farmer JD, Bui QM. *et al.* Statistical basis for predicting technological progress. *PloS ONE* 2013;**8**:e52669.
59. Farmer JD, Lafond F. How predictable is technological progress?. *Res Pol* 2016;**45**:647–65.
60. Panjer HH. *Operational Risk: Modeling Analytics*. Hoboken: John Wiley & Sons, 2006.
61. Mikosch T., Non-life insurance mathematics: an introduction with stochastic processesMikosch. Berlin: Springer, 2004.
62. Woods DW, Moore T, Simpson AC., The county fair cyber loss distribution: drawing inferences from insurance prices. *Digit Threats Res Pract* 2021;**2**:1–21.
63. Nadarajah S, Zhang Y, Pogány TK., On sums of independent generalized Pareto random variables with applications to insurance and CAT bonds. *Probab Eng Inform Sci* 2018;**32**:296–305.
64. Wolfram S. *Mathematica: A System for Doing Mathematics by Computer*. Boston: Addison Wesley Longman Publishing Co. Inc., 1991.

65. Panjer HH, Willmot GE. Insurance risk models. Schaumburg: Society of Actuaries, 1992.
66. Lin D, White JM, Byrne S,. *et al*. JuliaStats/Distributions.jl: a Julia package for probability distributions and associated functions. GitHub, 2019. https://doi.org/10.5281/zenodo.2647458.
67. Johnson SG., QuadGK.jl: Gauss–Kronrod integration in Julia. GitHub, 2013. https://github.com/JuliaMath/QuadGK.jl (18 December 2023, date last accessed).
68. Christ S, Schwabeneder D, Rackauckas C,. *et al*. Plots.jl – a user extendable plotting API for the julia programming language. *J OpenResearch Softw* 2023;**11**:15. https://openresearchsoftware.metajnl.com/articles/10.5334/jors.431/ (18 December 2023, date last accessed).
69. Benjamin S. Loadings for insurance premiums. *Geneva Pap Risk Insur* 1986;**11**:110–125.
70. Clark DR. Basics of reinsurance pricing?. CAS Actuarial Study Note. Arlington: Casualty Actuarial Society, 2014. https://www.casact.org/sites/default/files/old/studynotes_clark_2014.pdf (18 December 2023, date last accessed).
71. Lloyd's Minimum Standards MS5 - Risk Management. https://assets.lloyds.com/media/b9e65fca-1084-47e5-a057-25a855ede26d/MS5%20Risk%20Management.pdf (18 December 2023, date last accessed).
72. National Association of Insurance Commissioners. Report on the cyber insurance market. 2022. https://content.naic.org/sites/default/files/cmte-c-cyber-supplement-report-2022-for-data-year-2021.pdf (18 December 2023, date last accessed).
73. Cohn C. Insurers run from ransomware cover as losses mount. Reuters, 2021. https://www.reuters.com/markets/europe/insurers-run-ransomware-cover-losses-mount-2021-11-19/ (18 December 2023, date last accessed).
74. Mazzoccoli A, Naldi M. Robustness of optimal investment decisions in mixed insurance/investment cyber risk management. *Risk Anal* 2020;**40**:550–64.
75. Skeoch HR., Expanding the Gordon-Loeb model to cyber-insurance. *Comput Secur* 2022;**112**:102533.
76. Epanechnikov VA., Non-parametric estimation of a multivariate probability density. *Theor Probab Appl* 1969;**14**:153–8.
77. Woods DW., A turning point for cyber insurance. *Commun ACM* 2023;**66**:41–4.
78. Kasper D., Analyzing the feasibility of cyber bonds by stochastically solving a copula-based model with differential evolution. Ph.D. Thesis, Facultyof Management, Economics and Social Sciences, University ofCologne, 2019.
79. Braun A, Eling M, Jaenicke C., Cyber insurance-linked securities. *ASTIN Bull J IAA* 2023;**53**:1–22.
80. Kasper D, Grossklags J. A hierarchical macroeconomic copula model for cyber damages based on current cyber insurance prices. In: *International Conference on Science of Cyber Security*. Berlin: Springer, 2022, 472–83.
81. Bajoori E, Caulfield T, Ioannidis C., Cyber security service providers-should we leave them alone?. In: *Workshop on Approaches to Modelling Heterogeneous Interacting Systems. In Association with Financial Cryptography*. Grenada: International Financial Cryptography Association., 2022 .

Appendix 1: insurer loss distributions

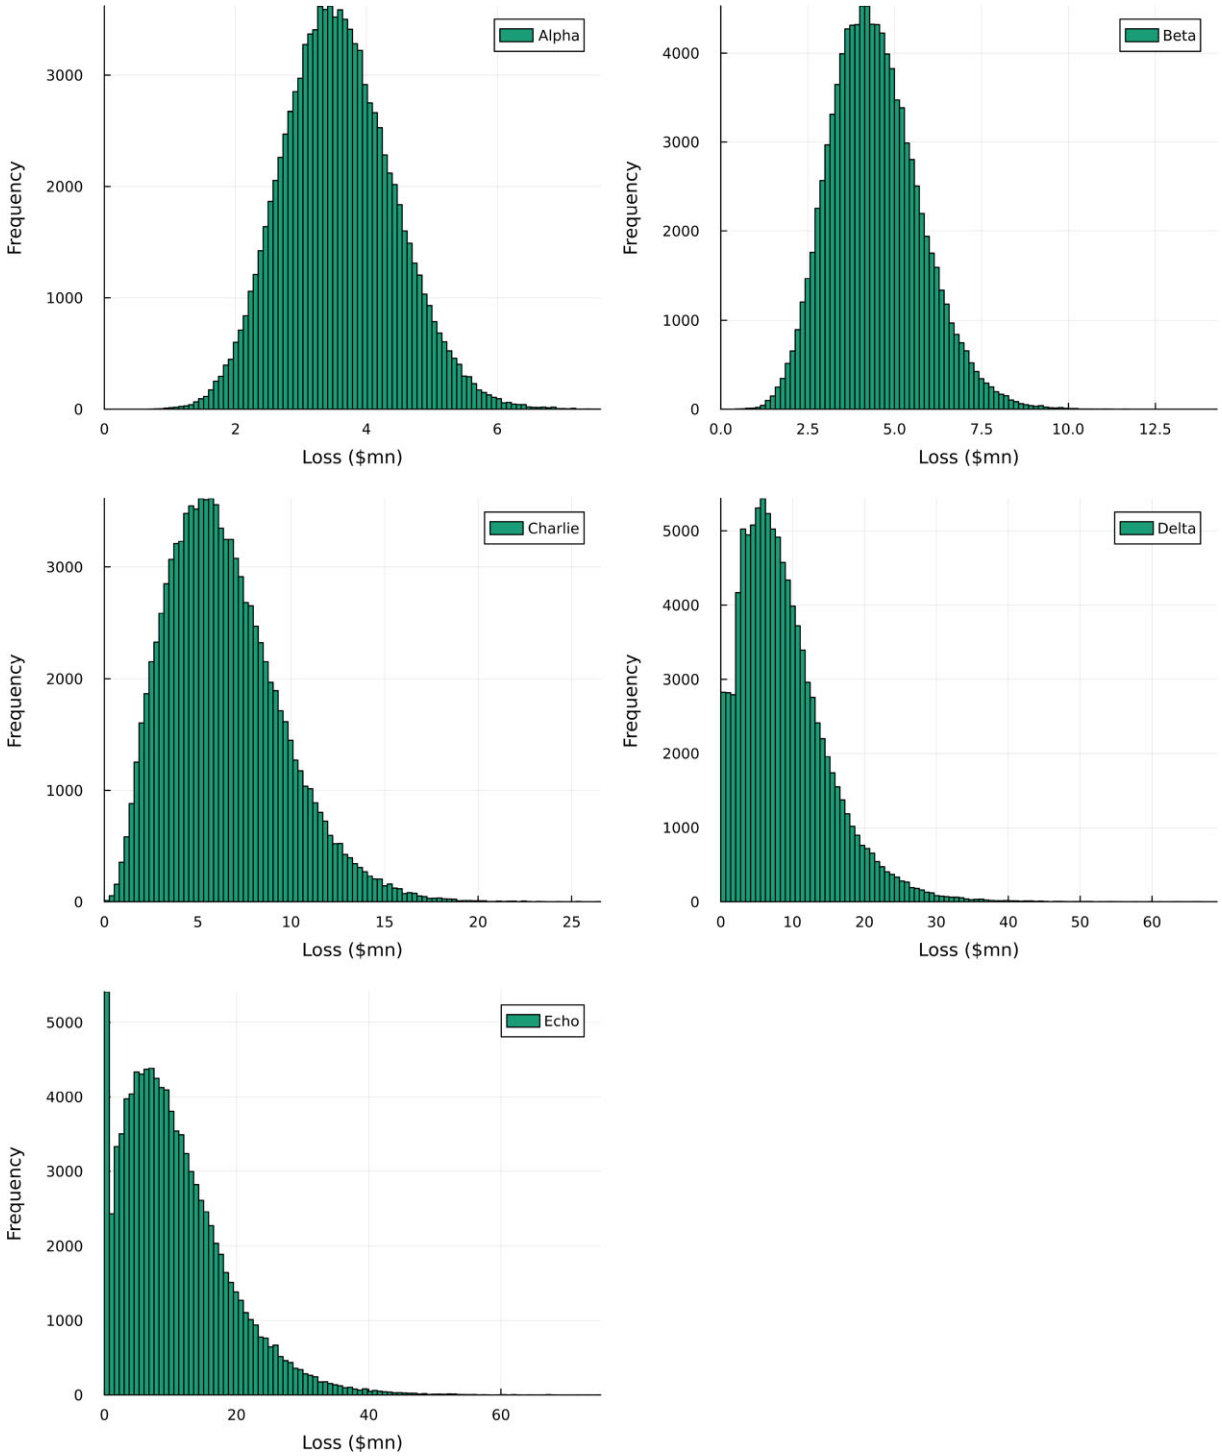

Figure A1. Insurer simulated loss distributions (Simulation procedure).