Voluntary Investment, Mandatory Minimums, or Cyber Insurance: What Minimizes Losses?

Adam Hastings
Columbia University

Simha Sethumadhavan
Columbia University

Abstract

In recent years there has been significant interest from policymakers in addressing ransomware through policy and regulations, yet this process remains far more of an art than a science. This paper introduces a novel method for quantitatively evaluating policy proposals: we create a simulated game theoretic agent-based economic model of security and use it as a testbed for several policy interventions, including a hands-off approach, mandatory minimum investments, and mandatory cyber insurance. Notably, we find that the bottleneck for better security outcomes lies not in better defender decision-making but in improved coordination between defenders: using our model, we find that a policy requiring defenders to invest at least 2% of resources into security each round produces better overall outcomes than leaving security investment decisions to defenders *even when the defenders are "perfect play" utility maximizers*. This provides evidence that security is a weakest-link game and makes the case for mandatory security minimums. Using our model, we also find that cyber insurance does little to improve overall outcomes. To make our tool accessible to others, we have made the code open source and released it as an online web application.

1 Introduction

Around the world, governments have signaled interest in improving security through policy, with proposals and regulations taking shape in the U.S., U.K., E.U., and beyond [1–5]. Motivating these proposals is the recognition that security is as much a problem of bad economics and misaligned incentives as it is one of insecure technology [6–8]; hence recent initiatives like the White House's National Cybersecurity Strategy have explicitly called for the re-shaping of market forces and the re-alignment of incentives to favor long-term investment [1]. However, while such high-level goals are promising, it remains unclear which policies are best suited to achieve these goals. Complicating matters is the "wicked problem" nature of cybersecurity, where policy interventions are costly and the efficacy of such interventions is difficult to predict *ex ante* or measure *ex post* [9, 10].

This work enables a quantitative approach towards security policymaking. In this paper, we develop an agent-based economic model of a cybersecurity ecosystem, with three classes of agents—defenders, attackers, and insurers—who exhibit game theoretic "perfect play", meaning they always choose the strategy that maximizes expected utility. However, our model is not strictly an analytic game theory model but instead is a hybrid approach that also uses Monte Carlo simulation and draws from empirically-derived inputs. To put this tool directly into the hands of others, we have released an interactive version of our model as a web application where users can adjust model inputs, run simulations, and produce model outputs themselves [11].

We then use the model to study new security policies and reaffirm existing ones. Using our model as if we were policymakers, we conduct several studies, summarized below:

Our first study explores voluntary security investments. In this hands-off approach, there is no explicit mandate or policy for security, and security spending is managed by the defenders themselves. Each round, defenders are free to invest any amount of their wealth into security, and are also free to purchase insurance policies from an insurer. Even under perfect play this model yields heavy losses for defenders who lose their wealth to ransom payments and recovery costs (later we discover that the poor result is not due to a flaw in defenders' strategies but due to a coordination problem). This study serves as a baseline model for further studies in this work.

Our second study is a sensitivity analysis. Our model provides the ability to observe how a change to an input (or inputs) influences the model output. This is especially relevant for policymakers, who—through various regulatory carrots and sticks—may have the ability to sway several of the real-world analogs of our model's inputs, and may use the model to observe the effects of such interventions. We perform a one-way sensitivity analysis on our model to observe which inputs the model output is most sensitive to. From this we find

that policymakers should try to reduce the number of attacks attempted and prevent the increase of recovery costs.

In our third study, we explore the effects of a policy where defenders must invest some minimum amount towards security. Surprisingly, we find that requiring defenders to invest 2% of resources into security at the beginning of each round leaves defenders better off than the baseline case where defenders make no mandatory investments but still make perfect utility-maximizing decisions.

Our fourth study concerns cyber insurance. For many organizations, cyber insurance has become an essential tool for managing risk. However, the role of insurance in reducing the incidence or impact of attacks is contentious [12–14]. In this study, we explore the effects of mandatory cyber insurance and find that its effects on outcomes are negligible.

In light of the previous study, one policy suggestion might be to lower premiums by means of a not-for-profit insurance scheme, e.g. akin to the U.S. Federal Deposit Insurance Corporation (FDIC) or the Federal Crop Insurance Corporation (FCIC). Hence our fifth study is a actuarially fair insurance scheme which makes insurance a more attractive option for defenders but does not improve overall outcomes compared to the baseline model. Finally, for our sixth study and for the sake of completeness, we also trial a version of our model where defenders' wealth grows over time.

The key contributions of this paper are as follows:

- We advance the state of the art in security economic modeling by combining the best features of game theory models, iterative games, empirical research, and large-scale Monte Carlo simulation to achieve a model of greater richness and detail than prior works.
- Using our model, we generate new insights about cybersecurity: We find that mandatory investments produce better outcomes than the uncoordinated every-defender-for-themselves scenario, even when defenders choose perfect play. From our model we also find that cyber insurance did not improve outcomes for defenders given the current state-of-affairs as captured via published ransomware data.
- We provide policymakers with a novel and interactive tool for exploring security tradeoffs and the effects of various policy interventions [11]. Notably, all the study presented in this paper can easily be reproduced simply by adjusting the input variables. In addition, we have made our simulator code open source so that other researchers may use and build upon our work [15, 16].

The paper is organized as follows: we situate our work within the context of prior scholarship in Section 2. In Section 3, we conduct empirical research to estimate real-world values and distributions of values for model input variables. Section 4 describes the model gameplay and agent strategies. Section 5 presents our baseline model (i.e. with no policy interventions). Section 6 presents a sensitivity analysis of the model, while Sections 7–10 use the model to explore the effects of mandatory security investments, mandatory cyber insurance, actuarially fair (i.e. profitless) insurance, and model behavior under compound growth, respectively. Section 11 describes the limitations of our work. This paper then concludes in Section 12.

2 Related Work

Our work is an economic model of security that builds upon several prior works in this area. However, we employ a novel technique by creating a hybridization of three categories of prior works: analytical models, empirical models, and simulation. To our knowledge, our work is the only economic security model to take such an approach.

To explain this difference, consider prior analytical models [17–27]. These works generally describe some aspect of security as a set of abstracted mathematical equations; then, using optimization methods or other analytic techniques, these works find critical points in the system (in game theoretic models, often in the form of Nash or Stackelberg equilibria). These works provide clean solutions but are often simplified down to only a handful of variables to enable optimization methods and often rely on parameters that are unknown in the real-world. Like the other analytic works, our work uses game theoretic decision-making but only to determine optimal agent strategies and not system equilibria. This allows us greater flexibility than what is possible with strictly analytic models.

We also draw inspiration from prior empirical works in security economics, which use (or collect) real-world datasets to model some element of security [28–31]. Likewise, in our own work, we collect real-world datasets of interest and then perform regressions and fit distributions to them; this is strictly different from the analytic works above. However, unlike other empirical works, our end goal is not to provide a description of the world but to use this description to create a more realistic data-driven model.

Our work also shares similarities with simulation-based economic models, which can allow for richer modeling than the strictly analytic models [32–34]. However, to our knowledge, our work is the only simulation-based model to investigate security investment decision-making which previously had only been attempted analytically [17, 19, 21, 22]. Furthermore, to our knowledge, prior simulation-based works try to discover optimal decision-making whereas our agents are *constructed* with optimal decision-making, owing to the analytic portions of our hybridized approach.

3 Model Setup

Our model contains three classes of players: defenders, attackers, and insurers. The defenders $\mathbf{D} = \{d_1, d_2, \ldots d_{|\mathbf{D}|^{[0]}}\}$ are the set of players who have assets they are trying to defend. The attackers are a set of agents $\mathbf{A} = \{a_1, a_2, \ldots a_{|\mathbf{A}|^{[0]}}\}$ that try to extract wealth from the defenders by means of ransomware. The insurers are the set of agents $\mathbf{I} = \{i_1, i_2, \ldots i_{N_I}\}$ who sell insurance policies to defenders. In this section, we use publicly-available data sources to determine appropriate values and value distributions for various model inputs.

3.1 Number of Agents

The first model input is the number of agents. An industry report finds that in the real world there are roughly 50 known active ransomware groups at any given time [35]. Hence we initialize our model with 50 attackers, i.e. $|\mathbf{A}| = 50$. In our notation, we also use brackets $[\,]$ to denote timestep, and so at model initialization (i.e. timestep $t = 0$) we write $|\mathbf{A}|^{[0]} = 50$.

From the same report we find that the 50 ransomware groups successfully attack roughly 5000 victims per year [35], or about 100 attacks per group, which we denote with $K = 100$. We can use K to determine the initial number of defenders $|\mathbf{D}|^{[0]}$: First, if there are $|\mathbf{A}|^{[0]} = 50$ attackers and $K = 100$ attacks per attacker per year, and $|\mathbf{D}|^{[0]}$ defenders, then the probability that a given defender d_i is paired with attacker a_i during timestep $t = 0$ is $K/|\mathbf{D}|^{[0]}$; assuming independence between attackers, the probability of a defender being attacked by *any* attacker is

$$\mathbb{P}\left[\text{attack}\right]^{[0]} = 1 - \left(1 - \frac{K}{|\mathbf{D}|^{[0]}}\right)^{|\mathbf{A}|^{[0]}} \tag{1}$$

A choice of $|\mathbf{D}|^{[0]} = 5000$ yields $\mathbb{P}\left[\text{attack}\right]^{[0]} = 0.636$, which is reasonably close to the percentage of organizations that have been hit with ransomware in the recent years [36].

To determine the initial number of insurers, we find that among the Fortune 500 companies, roughly 1 out of 25 are insurers who offer property and casualty insurance (including cyberinsurance). Creating one insurer for every 25 defenders gives $|\mathbf{I}|^{[0]} = |\mathbf{D}|^{[0]}/25 = 200$. Further validating this choice is that $|\mathbf{I}|^{[0]} = 200$ is within the same order of magnitude as the number of individual insurance companies offering cyberinsurance as reported to the National Association of Insurance Commissioners in 2022 [37].

3.2 Wealth Distribution

Each agent is initialized with some amount of wealth. To determine appropriate choices for agents' wealth, we first analyzed the revenue and earnings of the top global 1000 companies as ranked by market capitalization [38]. We find that revenue is lognormally distributed, and we fit a lognormal distribution to the data with $\mu_W = 1.135$ and $\sigma_W = 1.118$ (in terms of billions). A Kolmogorov-Smirnov test on the distribution yields $D = 0.0314$ and $p = 0.303$, indicating a good fit. We use this distribution to initialize our three classes of agents.

Defenders—at timestep $t = 0$, each defender d_i is initialized with wealth $d_{i,w}^{[0]}$ where

$$d_{i,w}^{[0]} = w, \quad w \sim \text{Lognormal}(\mu_w, \sigma_w^2) \tag{2}$$

Attackers—we rely on two assumptions to initialize attacker wealth: 1) attackers' wealth follows a similar lognormal distribution as defenders, and 2) attackers are on average less wealthy than the defenders they attack. To account for this we introduce an scaling factor s_I to represent the relative inequality between attackers and defenders. We estimate this parameter to be $s_I = 0.001$ based on available real-world data: we find that the largest ransomware organizations earn roughly \$100M a year in revenues [39–41]; in comparison, companies similar to those in our defender set have revenues on the order of \$100B a year [38]. Thus we find it reasonable to conclude that there is a 1:1000 ratio in size between attackers and defenders and set $s_I = 0.001$, and initialize attackers' wealth as follows:

$$a_{i,w}^{[0]} = s_I \cdot w, \quad w \sim \text{Lognormal}(\mu_w, \sigma_w^2) \tag{3}$$

Insurers—we initialize each insurer i_i's wealth $i_{i,w}^{[0]}$ by drawing from the same distribution as the defenders:

$$i_{i,w}^{[0]} = w, \quad w \sim \text{Lognormal}(\mu_w, \sigma_w^2) \tag{4}$$

3.3 Ransom Prices

The loss L of a ransomware attack is the sum of two components: the cost of the ransom itself and the cost of recovery from the attack (excluding the ransom). Prior work shows that both of these costs are largely a function of organization size [36], which we model using a defender d_i's wealth $d_{i,w}^{[t]}$:

$$L\left(d_{i,w}^{[t]}\right) = C_{\text{rans}}\left(d_{i,w}^{[t]}\right) + C_{\text{rec}}\left(d_{i,w}^{[t]}\right) \tag{5}$$

where $C_{\text{rans}}\left(d_{i,w}^{[t]}\right)$ is the cost of the ransom payment itself and $C_{\text{rec}}\left(d_{i,w}^{[t]}\right)$ is the recovery cost.

From available data we find that ransom sizes scale linearly with organization size [36], which we model with a linear regression:

$$C_{\text{rans}}\left(d_{i,w}^{[t]}\right) = \beta_{0,\text{rans}} + \beta_{1,\text{rans}} \cdot d_{i,w}^{[t]} \tag{6}$$

with $\beta_{0,\text{rans}} = 0.00121$ and $\beta_{1,\text{rans}} = 792145$ which achieves $R^2 = 0.921$, indicating a good fit.

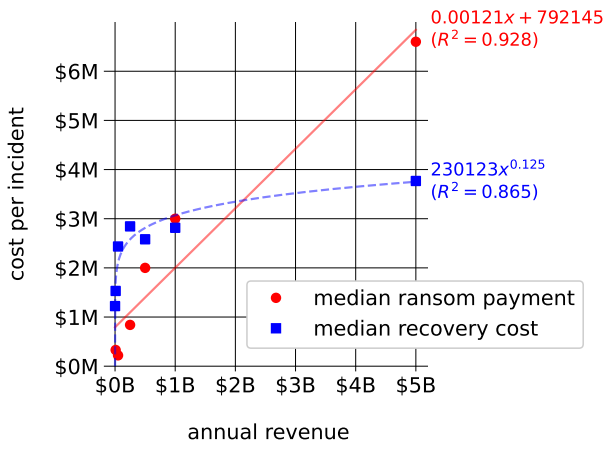

Figure 1: From an industry report we find that ransom payments and recovery costs scale linearly and sublinearly, respectively, with annual revenue [36]. Throughout this work, we use real-world data where possible to construct our model.

From the same dataset we find that recovery costs scale sublinearly, which we model as

$$C_{\text{rec}}\left(d_{i,w}^{[t]}\right) = a_{\text{rec}} \cdot \left(d_{i,w}^{[t]}\right)^{n_{\text{rec}}} \tag{7}$$

with $a_{\text{rec}} = 230123$ and $n_{\text{rec}} = 0.125$ which achieves $R^2 = 0.865$, indicating a satisfactory fit.

Hence in our model, ransom and recovery costs are scaled according to agent size and are given by Eqs. 6 and 7. Both are plotted in Fig. 1.

3.4 Security Posture

A key but nebulous value in real world security is security "posture", or the relative strength of an organization's security efforts. To estimate this distribution, we rely on a real-world proxy: the percentage of ransomware attacks that were successfully stopped before data was encrypted [36]. Following this definition, we represent defenders' security posture as the probability of preventing ransomware losses in the event of an attack. Fitting a normal distribution to this dataset yields mean $\mu_p = 0.28$ and standard deviation $\sigma_p = 0.10$. Hence at timestep $t = 0$, each defender d_i is initialized with posture $d_{i,p}^{[0]}$ where

$$d_{i,p}^{[0]} = p, \quad p \sim \mathcal{N}(\mu_p, \sigma_p^2), \quad 0 \le p \le 1 \tag{8}$$

A posture of $d_{i,p}^{[t]} = 0$ indicates that defender d_i has no security whatsoever (meaning any attack at all will be successful) while $d_{i,p}^{[t]} = 1$ indicates that defender d_i has perfect security (meaning no amount of attacking will be successful).

3.5 Attacker Profitability

Our model assumes that attackers must spend some amount of resources (the "wager") in order to successfully mount an attack. Before determining appropriate wager sizes (described later in §4), we must first establish ransomware agents' operating expense ratio $\text{OER} = \frac{\text{Total Expenses}}{\text{Total Revenue}}$. We establish parameters for this from two sources: First, we find that the Conti ransomware group had an estimated revenue of \$104.4M with \$31.2M in expenses (OER = 0.30) [42]. From another source, we find that Conti had an estimated return on investment (ROI) of +163% which implies OER = 0.38 [43]. From these sources we estimate that our attackers should make roughly \$3 for every \$1 they spend and choose $\text{OER} = \frac{1}{3}$.

Using our choice of OER, we can derive a formula for the "wager" an attacker must forfeit to attempt an attack. First, note that the expected revenue from an attack on defender d_i is a function of d_i's wealth and posture:

$$\mathbb{E}[\text{ revenue} \mid d_i] = \left(1 - \mathbb{E}\left[d_{i,p}^{[t]}\right]\right) \cdot \mathbb{E}\left[C_{\text{rans}}\left(d_{i,w}^{[t]}\right)\right] \tag{9}$$

Second, we define the expected expenses of an attack to scale linearly with d_i's posture and wealth, and subject to a scalar factor s_c:

$$\mathbb{E}[\text{expenses} \mid d_i] = s_c \cdot \mathbb{E}\left[d_{i,p}^{[t]}\right] \cdot \mathbb{E}\left[C_{\text{rans}}\left(d_{i,w}^{[t]}\right)\right] \tag{10}$$

Then

$$\text{OER} = \frac{\mathbb{E}[\text{expenses} \mid d_i]}{\mathbb{E}[\text{revenue} \mid d_i]} = \frac{s_c \cdot \mathbb{E}\left[d_{i,p}^{[t]}\right]}{\left(1 - \mathbb{E}\left[d_{i,p}^{[t]}\right]\right)} = \frac{1}{3} \tag{11}$$

Since $\mathbb{E}\left[d_{i,p}^{[0]}\right] = \mu_p = 0.28$ it follows that $s_c = 0.857$.

3.6 Security Investment Payoff

To determine the relationship between security investment and security posture, we build on assumptions from the Gordon-Loeb model [17]. Namely, we first borrow the assumption that a defender's security posture is a monotonic function of the amount of resources they have invested into security, and that security investments suffer from diminishing returns and asymptotically approach (but never reach) perfect security (i.e. $d_{i,p}^{[t]} = 1$). We also add the assumptions that zero security investment should produce zero security posture and that the effect of a security investment is relative to the size of the organization[1]. Finally, we assume that typical security spending should produce typical security posture.

[1] This last point also roughly follows from the Gordon-Loeb model [17] where investment z is relative to total possible asset loss λ. In our model, we assume λ to be the total assets of the defender rather than the value of a particular dataset. Another reason we choose to model investment as a percentage of wealth (instead of perhaps by modeling specific security controls) is to limit model complexity.

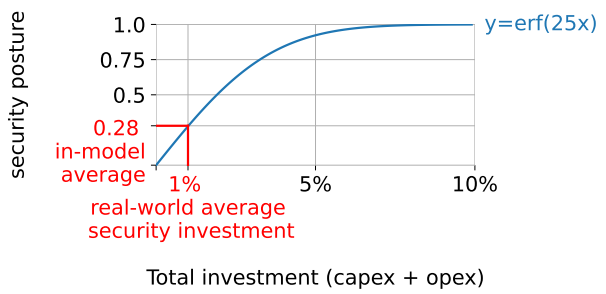

Figure 2: In our model, defenders can improve their security posture by making investments into security. We borrow assumptions on security investment from the Gordon-Loeb model [17], namely, that investments into security monotonically increase security posture but are subject to diminishing returns. We add the constraint that average security spending should produce average security posture. These constraints are satisfied by the function $y = \text{erf}(25x)$.

The above assumptions form a set of constraints. One possible class of functions that satisfy these constraints are sigmoidal functions. We choose the Gaussian error function erf as a candidate sigmoidal function and model it as a function of current *and previous* security spending (full details in §3.7). To satisfy the above final constraint (that average spending should produce average posture), we must scale the erf input by a scaling parameter s_p such that average investment produces average posture, i.e.

$$d^{[t]}_{i,\text{investment \%}} = \frac{\text{\$ amount of investment}}{d^{[t]}_{i,w}} \tag{12}$$

$$d^{[t]}_{i,\text{total investment \%}} = d^{[t]}_{i,\text{investment \%}} + d^{[t-1]}_{i,\text{investment \%}} \tag{13}$$

$$d^{[t]}_{i,p} = \text{erf}\left(s_p \cdot d^{[t]}_{i,\text{total investment \%}}\right) \tag{14}$$

We find that average security investment in the real world is about 1% of resources per annum [44, 45]. Recall that we empirically find average posture to be $\mu_p = 0.28$. Putting this all together, we require that

$$0.28 = \text{erf}(s_p \cdot 0.01) \tag{15}$$

which is satisfied when $s_p = 25$, and shown in Fig. 2.

3.7 Security Depreciation

Unfortunately, security posture depreciates without continued investment. We use a few assumptions to create a reasonable security depreciation schedule for defenders' security posture. First, we note that security investments can be in the form of *capital expenditures* ("capex") or *operational expenditures* ("opex")[2].

Second, we find that in most organizations security spending is dominated by operational expenditures with capital expenditures only consisting about 33% of security spending [46]. Hence we define $s_K = 1/3$ to be the percentage of security spending that retains value in future rounds (capex) while the other two-thirds retain no value in future rounds (opex). We also assume that capex decays by some fixed depreciation scalar λ where $0 \leq \lambda \leq 1$. Together,

$$d^{[t]}_{i,\text{opex}} = \frac{2}{3} \cdot d^{[t]}_{i,\text{investment \%}} \tag{16}$$

$$d^{[t]}_{i,\text{capex}} = \frac{1}{3} \cdot d^{[t]}_{i,\text{investment \%}} + (1-\lambda) \cdot d^{[t-1]}_{i,\text{capex}} \tag{17}$$

$$d^{[t]}_{i,p} = \text{erf}\left(s_p \cdot \left(d^{[t]}_{i,\text{capex}} + d^{[t]}_{i,\text{opex}}\right)\right) \tag{18}$$

To determine the rate of capex decay λ we again rely on the assumption that average security spending should produce average security posture. Namely, even with capex decay it should still be such that a 1% investment of resources produces an average security posture $\mu_p = 0.28$ *in perpetuity*. By simulation we find that this constraint is satisfied when $\lambda = 0.4$, meaning that the value that capital expenditures provide towards security posture depreciates by 40% every iteration, and that even with this depreciation an organization investing 1% of resources into security each iteration will maintain a security posture of $\mu_p = 0.28$.

3.8 Insurance Policy Parameters

From a dataset of real-world cyberinsurance policies we find that premiums have a linear relationship with retentions [30]. Hence in our model, insurers will sell a policy Π with premium Π_P and retention Π_R where $\Pi_R = 25\Pi_P$ ($R^2 = 0.90$). Premium prices themselves are calculated during runtime and are based on in-model risk (see §4.3.3).

We also include a loss ratio parameter: as part of our model, insurers operate under a fixed target loss ratio *LR*, which is defined as the percentage of collected premiums that are paid to policyholders in the forms of claims. We take inspiration from the United States' Affordable Care Act which mandates a loss ratio of 80% and hence we choose $LR = 0.80$.

4 Model Gameplay

Once game initialization is complete, gameplay begins. Our model is an **iterated** system that evolves over a series of round timesteps $t = [1, 2, \ldots, M]$. There are five major steps that happen in each round:

[2] A capital expenditure is a purchase that retains its value upon purchase, such as a hardware firewall device. An operational expenditure is one that is "consumed" upon purchase, such as paying for services like incident response.

4.1 Security Depreciation

For all rounds $t > 0$, defenders' prior capital expenditures experience decay given by

$$d_{i,\text{capex}}^{[t]} = d_{i,\text{capex}}^{[t-1]} \cdot (1-\lambda) \tag{19}$$

4.2 Threat Analyses

At the start of each round, each player performs a "threat analysis" to evaluate current risk levels to inform decision-making later in the round.

4.2.1 Attacker Threat Analysis

Attackers are only incentivized to attack if the expected gains from attacking a victim exceed the expected loss, which is partially dependent on the victim's security posture. However, attackers do not know victims' posture prior to attacking and instead must rely on an expected posture. Hence attackers compute a method-of-moments estimation of the average defender posture $\widehat{\mu_p}^{[t]}$ [47].

Recall from §3.5 our definitions of expected attacker gains and losses. During gameplay, attackers use their approximation $\widehat{\mu_p}^{[t]}$ in lieu of $\mathbb{E}\left[d_{i,p}^{[t]}\right]$, reducing to the following:

$$\mathbb{E}\left[\text{attacker earnings} \mid d_i^{[t]}\right] = C_{\text{rans}}\left(d_{i,w}^{[t]}\right) \cdot \left(1-\widehat{\mu_p}^{[t]}\right) \tag{20}$$

$$\mathbb{E}\left[\text{attacker loss} \mid d_i^{[t]}\right] = s_p \cdot \widehat{\mu_p}^{[t]} \cdot C_{\text{rans}}\left(d_{i,w}^{[t]}\right) \tag{21}$$

Then, given a random pairing between attacker $a_i^{[t]}$ and defender $d_i^{[t]}$, the attacker will only attempt an attack if

$$\mathbb{E}\left[\text{attacker earnings} \mid d_{i,w}^{[t]}\right] > \mathbb{E}\left[\text{attacker loss} \mid d_{i,w}^{[t]}\right] \tag{22}$$

4.2.2 Insurer Threat Analysis

Insurers must determine the expected number of attacks that a given defender d_i will face during round t in order to accurately price policies they will sell later in the round. Recall that each attacker will attack K victims at random each round. The probability that a given defender d_i will be attacked by a given attacker a_j during round t is $K/|\mathbf{D}|^{[t]}$ and the probability of a defender $d_i^{[t]}$ being attacked by *any* attacker during round t is

$$\mathbb{P}\left[\text{attack}\right]^{[t]} = 1 - \left(1 - \frac{K}{|\mathbf{D}|^{[t]}}\right)^{|\mathbf{A}|^{[t]}} \tag{23}$$

We make the simplifying assumption that each defender can be attacked only once per round. We also assume that insurers are able to accurately determine the number of active attackers $|\mathbf{A}|^{[t]}$ as part of their operating expenses. The value of $|\mathbf{A}|^{[t]}$ is considered business intelligence and is not disclosed to the defenders.

To further improve risk estimates, insurers are also able to model attackers' behavior. This first includes re-calculating the attacker's estimate of defender's average posture $\widehat{\mu_p}^{[t]}$. From this, insurers are able to determine the attacker's expectations for attacking, namely Eq. 22.

Second, insurers can also determine the probability that a given attacker has enough assets to even attempt an attack on the average defender. To this end, insurers compute a method-of-moments estimation of the distribution parameters of attackers' wealth with mean $\widehat{\mu_A}^{[t]}$ and standard deviation $\widehat{\sigma_A}^{[t]}$ [47]. Then given a random pairing between some defender $d_i^{[t]}$ and attacker $a_i^{[t]}$ the probability that the attacker has enough funds to attack is given by

$$\mathbb{P}\left[\text{attempt}|d_{i,w}^{[t]}\right] = \mathbb{P}\left[a_{i,w}^{[t]} \geq \mathbb{E}\left[\text{attacker loss} \mid d_{i,w}^{[t]}\right]\right] \tag{24}$$

$$= \Phi\left(\frac{\ln d_{i,w}^{[t]} - \widehat{\sigma_A}^{[t]}}{\widehat{\mu_A}^{[t]}}\right) \tag{25}$$

where Φ is the CDF of the standard normal distribution.

Using this threat analysis, insurers estimate the probability of a given defender experiencing a loss during round t:

$$\mathbb{P}\left[\text{defender loss} \mid d_i^{[t]}\right] = \mathbb{P}\left[\text{attack}\right]^{[t]} \cdot \mathbb{P}\left[\text{attempt} \mid d_{i,w}^{[t]}\right] \cdot \left(1 - d_{i,p}^{[t]}\right) \tag{26}$$

4.2.3 Defender Threat Analysis

Each defender also computes the probability of an attack. However, defenders do not have access to the same actuarial information as the insurers and rely on an approximation instead. To construct this approximation, we use two heuristics: first we assume that defenders know the percentage of fellow defenders that are looted each round but not necessarily the baseline number of attacks:

$$\widehat{\mathbb{P}}\left[\text{loss}\right] = \frac{\#\text{ defenders ransomed during round } (t-1)}{|\mathbf{D}|^{[t-1]}} \tag{27}$$

Second, we assume that defenders know their own security posture but not the postures of others, and rely on the assumption that other defenders share their same posture, i.e. $\widehat{\mu_p} = d_{i,p}^{[t]}$. Hence defenders are able to work backward to derive an estimated baseline rate of attack:

$$\widehat{\mathbb{P}}\left[\text{attack}\right]^{[t]} = \frac{\widehat{\mathbb{P}}\left[\text{loss}\right]}{1-\mu_p} \tag{28}$$

Using this rudimentary threat analysis, defenders compute the probability of a loss in the upcoming round as

$$\widehat{\mathbb{P}}\left[\text{loss} \mid d_{i,p}^{[t]}\right] = \widehat{\mathbb{P}}\left[\text{attack}\right]^{[t]} \cdot \left(1 - d_{i,p}^{[t]}\right) \tag{29}$$

4.3 Defender Strategy Selection

After threat analysis, defenders choose a strategy for the upcoming round. Each defender d_i has three options available: invest in security, buy an insurance policy, or do nothing.

4.3.1 Strategy I: Do Nothing

If a defender chooses to neither invest in security nor buy insurance, the expected loss is

$$\mathbb{E}\left[\text{loss} \mid d_i^{[t]}\right] = \widehat{\mathbb{P}}\left[\text{loss} \mid d_{i,p}^{[t]}\right] \cdot L\left(d_{i,w}^{[t]}\right) \tag{30}$$

which follows from Eqs. 5 and 29.

4.3.2 Strategy II: Invest in Security

Each defender may also choose to invest some amount $0 \leq x \leq d_{i,w}^{[t]}$ into security. From §3 we can write posture as a function of investment x:

$$d_{i,p}^{[t]}(x) = \text{erf}\left(s_c \cdot \left(x + d_{i,\text{capex}}^{[t-1]}\right) / d_{i,w}^{[t]}\right) \tag{31}$$

The probability of ransom given investment x therefore is

$$\mathbb{P}\left[\text{loss} \mid d_i^{[t]}, x\right] = \widehat{\mathbb{P}}[\text{attack}]^{[t]} \cdot \left(1 - d_{i,p}^{[t]}(x)\right) \tag{32}$$

and the expected loss is

$$\mathbb{E}\left[\text{loss} \mid d_i^{[t]}, x\right] = \mathbb{P}\left[\text{loss} \mid d_i^{[t]}, x\right] \cdot L\left(d_{i,w}^{[t]} - x\right) + x \tag{33}$$

which is a convex function. Defenders then use Brent method [48] to find the optimal investment amount x^* that minimizes Eq. 33. Hence the expected loss with optimal investment x^* is

$$\mathbb{E}\left[\text{loss} \mid d_i^{[t]}, x^*\right] = \mathbb{P}\left[\text{loss} \mid d_i^{[t]}, x^*\right] \cdot L\left(d_{i,w}^{[t]} - x^*\right) + x^* \tag{34}$$

4.3.3 Strategy III: Buy Insurance

The third option is to buy an insurance policy. We allow each defender to request $Q = 10$ quotes from the insurers, chosen at random. As part of requesting a quote, the insurer conducts an "audit" of the defender to obtain $d_{i,w}^{[t]}$ and $\widehat{d}_{i,p}^{[t]}$, an estimate of d_i's posture. We give insurers an *estimate* of d_i's posture to model the information asymmetry between insurers and insurees and the difficulty of measuring security posture that is inherent to the underwriting process [30]. The insurer calculates the probability of d_i experiencing a loss during the policy given $d_{i,w}^{[t]}$ and $\widehat{d}_{i,p}^{[t]}$:

$$\begin{aligned}\mathbb{P}\left[\text{loss} \mid d_{i,w}^{[t]}, \widehat{d}_{i,p}^{[t]}\right] = \\ \mathbb{P}[\text{attack}]^{[t]} \cdot \mathbb{P}\left[\text{attempt} \mid d_{i,w}^{[t]}\right] \cdot \left(1 - \widehat{d}_{i,p}^{[t]}\right)\end{aligned} \tag{35}$$

Insurers then write a policy Π with premium Π_P and retention Π_R that is expected to achieve the target loss ratio LR, i.e.

$$LR = \frac{\mathbb{E}\left[\text{insurer loss} \mid \Pi\right]}{\mathbb{E}\left[\text{insurer gain} \mid \Pi\right]} \tag{36}$$

$$= \frac{\mathbb{P}\left[\text{loss} \mid d_{i,p}^{[t]}\right]\left(L\left(d_{i,w}^{[t]}\right) - \Pi_R\right)}{\Pi_P} \tag{37}$$

Recall from §3 that $\Pi_R = 25 \cdot \Pi_P$. Solving for Π_P yields

$$\Pi_P = \frac{\mathbb{P}\left[\text{loss} \mid d_{i,p}^{[t]}\right] \cdot L\left(d_{i,w}^{[t]}\right)}{LR + 25 \cdot \mathbb{P}\left[\text{loss} \mid d_{i,p}^{[t]}\right]} \tag{38}$$

and the policy $\Pi = \{\Pi_P, \Pi_R\}$ is given as a quote to the defender. Defenders then compute the expected loss given Π:

$$\mathbb{E}\left[\text{defender loss} \mid \Pi\right] = \Pi_P + \Pi_R \cdot \widehat{\mathbb{P}}\left[\text{loss} \mid d_{i,p}^{[t]}\right] \tag{39}$$

If the defender chooses to buy insurance, the insurer will buy a policy from the insurer that offered the lowest premium.

Choosing a strategy—the defender will choose the strategy that minimizes the expected loss as given by Eqs. 30, 33, and 39.

4.4 Fight

Once defenders have chosen their strategies for the round, the attacks begin. Each attacker targets K defenders, chosen at random.

Once paired with a defender d_i, an attacker a_i will compute the expected gains and losses to determine whether attempting an attack on d_i is financially rational (Eq. 22). If so, and if the attacker has enough wealth (Eq. 24), a_i will spend the amount given by Eq. 10 and attempt an attack.

With probability $d_{i,p}^{[t]}$ the defender "wins" the fight (meaning their security posture was sufficient to ward off the attack), and with probability $1 - d_{i,p}^{[t]}$, the attacker wins the fight. If the attacker wins, the defender d_i pays the ransom $C_{\text{rans}}\left(d_{i,w}^{[t]}\right)$ to attacker a_i. Defender d_i also loses wealth in the form of recovery costs $C_{\text{rec}}\left(d_{i,w}^{[t]}\right)$. If the defender had chosen insurance during Step 3, they are covered for their losses minus the retention Π_R.

4.5 Iterate Until Convergence

The above steps are iterated until the game reaches one of three termination conditions: Either all defenders die off ($d_{i,w} = 0 \quad \forall d_i \in \mathbf{D}$), all attackers die off ($a_{i,w} = 0 \quad \forall a_i \in \mathbf{A}$), or the game reaches a stable equilibrium. To define this last condition we consider a game stabilized if there are no attacks attempted for some number $\delta = |\mathbf{A}|^{[0]} = 50$ rounds in a row.

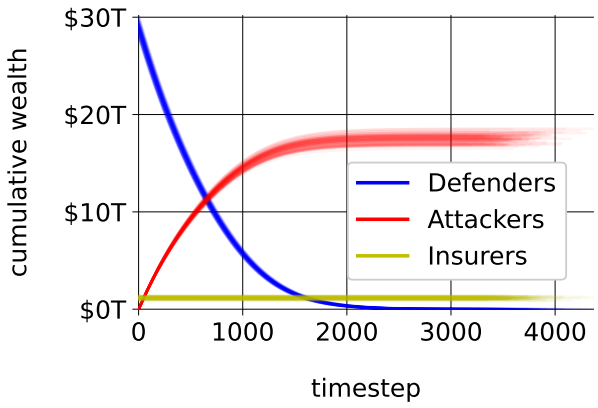

Figure 3: This figure plots the results of 100 simulations of the baseline model. In all 100 simulations, the defenders end gameplay with $0, with much their wealth transferred to the attackers via ransom payments (the rest having been lost to recovery costs, insurance premiums, or security investments). While hardly evident here, insurers are collecting premiums and paying claims throughout gameplay but their wealth remains largely unchanged.

5 Baseline Model Behavior

Our first study is to establish model behavior in the absence of any policy intervention, i.e. the model behavior when defenders are free to make utility-maximizing investments in security and insurance.

Study Configuration: Model inputs are subject to the default values as found in §3. To capture the range of possible model behavior, we ran the model 100 times.

Study Results: We highlight four results to summarize the baseline model: First, we observe that in the baseline model the attackers completely plunder the defenders (Figure 3).

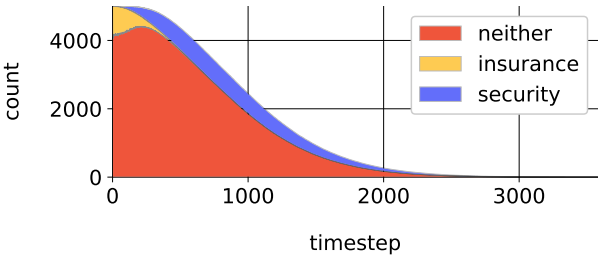

Figure 4: Stacked time series showing defenders' choices given the model's default input values. Results are averaged across 100 runs. Throughout simulation, the dominant strategy is to neither buy insurance nor invest in security. However, purchasing security becomes a more popular strategy as gameplay progresses (as a percentage of defenders).

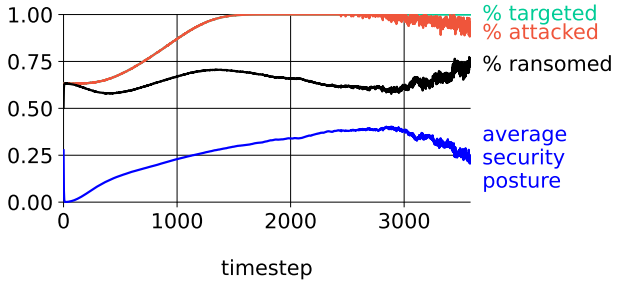

Figure 5: To explicate model behavior, we track the values of several "barometer" variables throughout model simulation. This figure shows four such variables averaged across 100 runs of the baseline model. We find that as the simulation progresses, defenders are ransomed and die off, causing the probability of any defender getting targeted and attacked in a given round to approach 1. We also find that average security posture starts at 0.28 (by construction—see §3) but quickly drops to near-zero due to a lack of security investments during the early stages of gameplay (compare with defender choices in Fig. 4). Later, incentives for defenders change and average security posture begins to rise but this is still not enough to prevent a total loss for defenders (compare with Fig. 3).

Second, we observe that defenders sometimes choose to invest in security and sometimes purchase insurance but in most cases are choosing to do neither (Figure 4).

Third, we observe interesting gameplay dynamics, as shown in Figure 5. Specifically, we find that average defender posture starts at $\mu_p = 0.28$ (which follows from Eq. 8); however, average posture quickly craters to near-zero due to the model's built-in security depreciation schedule (Eq. 18) and the lack of security investments early during gameplay (Fig. 4). After this initial drop, security investments increase and posture rises until around timestep 3000 where the game starts to equilibrate. Other gameplay conditions exhibit interesting dynamics as well: as defenders die off, the probability of getting targeted by an attacker approaches 1, and so does the probability of an attacker rationally deciding to attack (Eq. 22) (although this falls as the game equilibrates). However, the probability of a defender getting attacked and the attack succeeding (% ransomed) fluctuates considerably.

Finally, our fourth observation is that a considerable amount of wealth is lost to recovery costs (Fig. 6). From Figure 6 we also note that attackers are able to ransom considerable wealth (∼$17T) while only spending a fraction of this on expenditures (∼$300B) despite our construction of the wager (Eq. 10). This is explained by the low average security postures in the first ∼1000 rounds of execution (Fig. 5) where the bulk of ransom is stolen (Fig. 3).

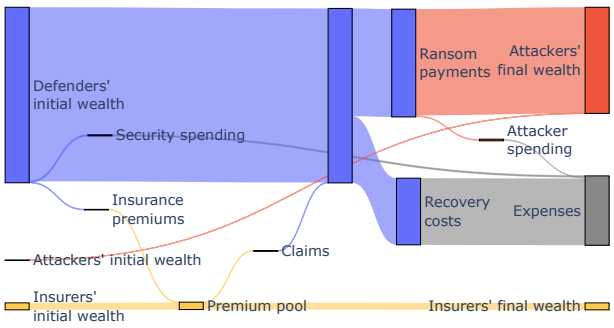

Figure 6: A Sankey diagram showing the transfer of wealth between the three classes of agents in the baseline model, averaged across 100 simulations. The majority of defender wealth is lost to ransom payments and recovery costs. In the baseline model, comparatively little is spent on security investments or insurance.

5.1 Model Validation

How do we know that our model is a reasonable abstraction of the real world? One method is to confirm that assumptions about the state of the world in §3 are reflected during simulation, or at least at the start of simulation. To this end, we identify several "canary" variables to confirm that various conditions are met at timestep $t = 0$. For example, we find that average initial defender posture $\bar{\mu_p}$ matches the distribution of known real-world security posture $\mu_p = 0.28$. Likewise, the initial probability of attack $\mathbb{P}[\text{attack}]^{[0]} = 0.59$ matches real-world rate of attacks [36]. Finally, we validate that defender strategies mimic the real world, with defenders choosing mixed strategies of buying insurance, investing in security, and doing neither.

However, we do not claim strict ecological validity of our model: clearly we do not live in world where defenders are completely looted of their wealth. Nevertheless, when there was a design tradeoff between ecological validity and model simplicity/interpretability, we chose the latter over the former. For example, see Section 10 where we improve validity, albeit (in our view) at the expense of model interpretability.

Besides model validation, there is also the concern that our C++ code is a correct implementation of the model given in Sections 3 and 4. To this end, we add runtime assertions to each variable (for example, asserting that variables representing probabilities are always between 0 and 1); we also add considerable "bookkeeping" assertions to the code to ensure that certain constraints are always satisfied (for example, that Σ initial wealth $=$ Σ end wealth $+$ Σ expenses).

6 Sensitivity Analysis

Given that we do not claim strict ecological validity, the value of this work lies not in the output values themselves but in the shapes of the outputs and how the outputs respond to inputs. For example, one may wish to understand which inputs have the greatest effect on outputs. This may be of particular interest to policymakers, who may have some degree of control over one or more of the real-world analogs to model inputs and naturally would be interested in identifying which of these input values, if adjusted, would provide the greatest benefit to the system at the least cost. This question can be addressed by means of our second study, a one-way sensitivity analysis of our model.

Study Configuration: To perform a sensitivity analysis, we sweep inputs across a range of plausible inputs to determine how variations in single inputs affect simulation outcomes. This requires two preliminary steps:

First, we must define a loss function to track the "goodness" of the outcome, which we define as the percentage of defender wealth remaining at timestep t:

$$\text{loss}[t] = \frac{\sum_{i=0}^{|\mathbf{D}|^{[0]}} d_{i,w}^{[0]} - \sum_{i=0}^{|\mathbf{D}|^{[t]}} d_{i,w}^{[t]}}{\sum_{i=0}^{|\mathbf{D}|^{[0]}} d_{i,w}^{[0]}} \tag{40}$$

Second, we must determine a reference timestep t at which to analyze system loss. While the obvious choice might be to evaluate loss at simulation termination (defined in §4.5), we find that perturbations of individual input values are generally not enough to forestall total defender loss (i.e. loss=100%) regardless of input value and ultimately obscure how inputs affect outcomes. Instead, we choose to perform our sensitivity analysis at some point during simulation where input value perturbations can have a visible affect on loss, which occurs when defenders wealth has been reduced halfway. In the baseline model, we see from Fig. 3 that this occurs at approximately $t = 500$[3].

Study Results: Sensitivity analysis results are given in Fig. 7. For model inputs that are expressed in terms of a percentage, we evaluate the model across the entire possible range of values $[0, 1]$ (Fig. 7a). Other model inputs are not bound by a range, so we evaluate on the range from $0\times$ to $2\times$ the baseline default values.We highlight three observations:

First, we observe that a number of variables are positively correlated with system loss but in several different ways. For example, the coefficient $\beta_{1,\text{ransom}}$ exhibits a near-linear effect on loss; this follows since $\beta_{1,ransom}$ is a linear scaling factor for computing ransom prices). In contrast, we see that K (the number of attacks an attacker attempts each round) has a sub-linear effect on loss, likely because the number of

[3]While not shown here, we perform sensitivity analyses at other timesteps to confirm that $t = 500$ is both in line with sensitivity analyses at other timesteps and also the most illustrative and informative timestep to perform a sensitivity analysis.

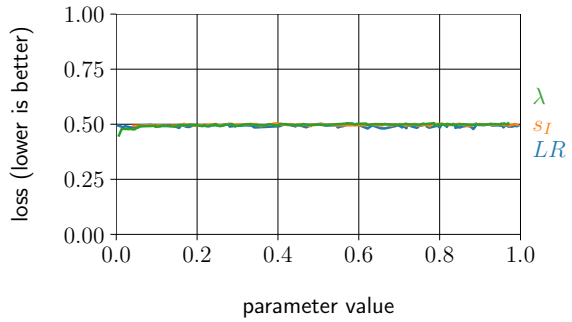

(a) Sensitivity analysis of model inputs that are expressed in terms of a percentage. For these inputs, we evaluate the model across the full range of possible input values (namely from 0 to 1).

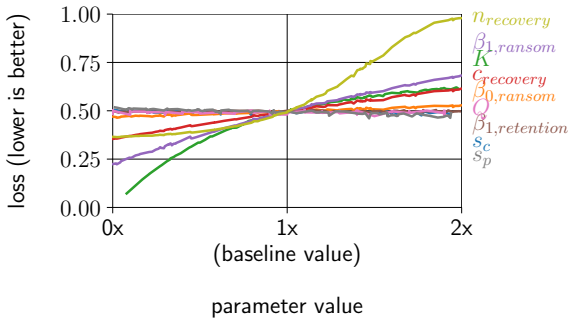

(b) Sensitivity analysis of non-percentage model inputs. For these inputs, we evaluate from 0× to 2× the default value.

Figure 7: Results from the model input sensitivity analysis, evaluated at timestep $t = 500$.

available victims saturates as attackers begin to compete for the same victims. Yet another: The exponent to the recovery cost function n_{recovery} exhibits a sigmoidal effect on loss. We also observe that some variables have a slightly negative correlation with loss. For example, increasing the security investment scaling factor s_p allows defenders to more efficiently translate security investments into elevated security posture, which ultimately reduces loss.

Second, we observe that several variables exhibit a "flat" sensitivity response, which *prima facie* might suggest the variables have no effect on system output. However, we stress that this is not necessarily so: our construction of loss is not the only possible metric of goodness, and other loss constructions (e.g. rate of loss, or loss including insurers' wealth) might show a different response. Another possibility is that the flat variables do not produce a significant effect at our given choice of timestep $t = 500$. As proof of such possibilities, consider §9, where we evaluate the effect of insurers who operate with loss ratio $LR = 100\%$: we observe multiple effects on gameplay despite what Fig. 7a may suggest.

Third, we can use the sensitivity analysis to help determine the weighted importance of each variable in the model. This is particularly useful for purposes of security policy and regulation, where one may attempt to modify the real-world system of security by influencing various parameters. For example, our sensitivity analysis shows that the exponent to the recovery cost function n_{recovery} produces catastrophic losses if allowed to increase; a regulator might want to investigate mechanisms for preventing the real-world analog of n_{recovery} from increasing. Interesting, the inverse is not true: If a regulator were to focus on trying to reduce loss by *decreasing* certain values, they should focus on decreasing the real-world analogs of our model variables K and $\beta_{1,\text{ransom}}$.

Finally, we point out that this is only a one-way sensitivity analysis involving the perturbation of a single input variable at a time. Nevertheless, policymakers—who may have influence over multiple real-world analogs of model inputs—may be interested in seeing the effects of adjusting two or more model inputs at a time. We do not include any multiple-way sensitivity analyses in this work (mostly due to the curse of dimensionality) but refer those who are interested in such analyses to use our online demo or open source code.

7 Mandated Security Investments

Amongst policymakers there has been discussion on how to stop "passing the buck" on cybersecurity by raising the standard for security [49]. In this study, we add a "policy" to the model, requiring defenders to invest some minimum percent of resources towards security to investigate how such a policy might improve (or hurt) overall outcomes for defenders. By extending the baseline model of Section 5, one may explore the effects of such a policy.

Study Configuration: To create this variant of our model, we first add a new parameter M, which is the percentage of resources a defender must invest in security at the start of each round of gameplay. As with voluntary security investments, the effect that investments have on security posture is given by Eq. 14. After the mandatory investment, defenders are still allowed to make a rational decision between insurance, additional security investment, or doing nothing.

Study Results: We evaluate our model at investment levels $M = \{1\%, 2\%, 3\%, 4\%, 5\%\}$. When $M = 1\%$, we find that the outcome is still a total loss for the defenders. However, the wealth is not being transferred to the attackers; instead, we observe that the bulk of defenders' wealth is now being spent on security investments instead (Fig. 8a). Does this mean the defenders are spending too much on security?

Remarkably, we find the answer to be **no**. Observe what happens when we increase M from 1% to 2%: unlike the cases

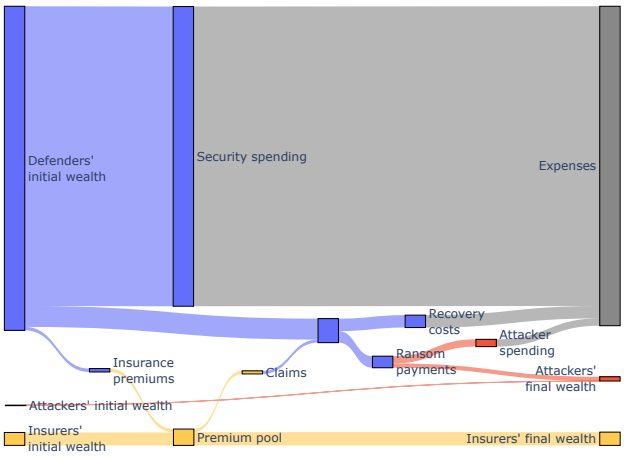

(a) Mandatory investment $M = 1\%$

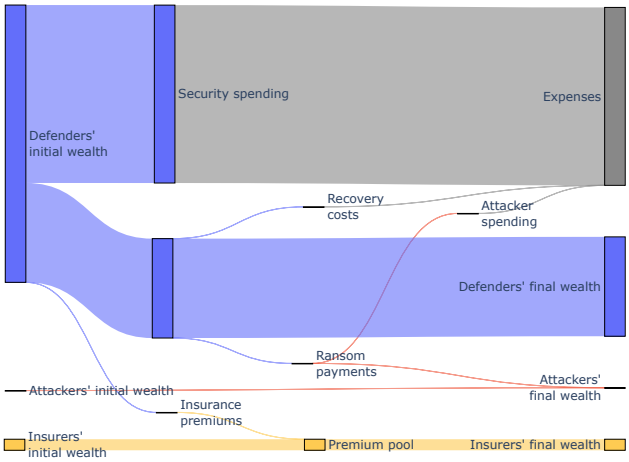

(b) Mandatory investment $M = 2\%$

Figure 8: Wealth transfers given a mandatory security investment M, averaged across 100 runs. Compare to the baseline case where $M = 0\%$ (Fig. 6). With a 1% mandatory minimum investment, the majority of defender wealth is spent on security spending (Fig. 8a). However, we also find that when the mandatory minimum investment is raised to 2%, defenders spend *less* on security overall and retain *more* wealth (Fig. 8b). The explanation for this counterintuitive result can be found in Fig. 9.

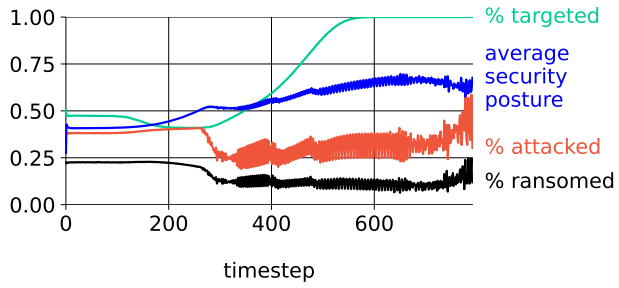

(a) Mandatory investment $M = 1\%$

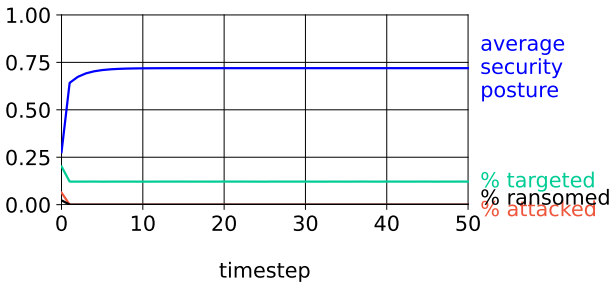

(b) Mandatory investment $M = 2\%$

Figure 9: Gameplay barometer values given a mandatory security investment M, averaged across 100 runs. Compare to the baseline case where $M = 0\%$ (Fig. 5). With a 1% mandatory minimum investment, average security posture remains higher than in the baseline case (never dropping below the initial 0.28), which reduces but does not eliminate the percentage of defenders that are attacked and ransomed each round. With a 2% mandatory investment, average security posture remains high enough to ward off all attacks altogether.

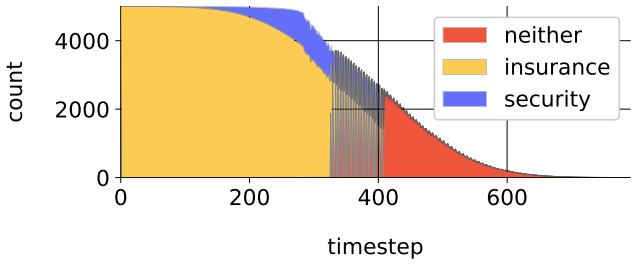

(a) Mandatory investment $M = 1\%$

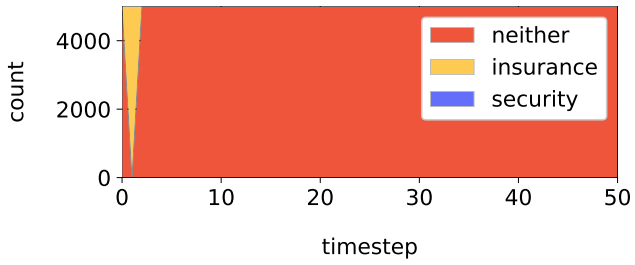

(b) Mandatory investment $M = 2\%$

Figure 10: Stacked time series showing defenders' choices given a mandatory security investment M, averaged across 100 runs. Compare to the baseline case where $M = 0\%$ (see Fig. 4). With a 1% mandatory investment, defenders initially favor insurance but also still make additional voluntary security investments. With a 2% mandatory investment, defenders find their security posture to be adequate and do not seek out additional security investments.

where $M = 0\%$ or $M = 1\%$, the defenders retain a sizeable portion of their initial wealth (Fig. 8b). Even more remarkably, the defenders managed to retain more wealth *while investing less in security!*

What causes this? Consider Fig. 9a: at a 1% mandate, average defender posture climbs to above 0.70 (significantly higher than the 0% mandate case in Fig. 5) and the % ransomed is considerably lower but the incentive to attack is never eliminated. We find that the attackers are able to prey on low-value targets (small assets and weak posture) long enough to stay alive; after ~ 700 rounds, even the strong-postured defenders have exhausted their wealth on security investments and are then killed off.

Now consider Fig. 9b: at a 2% mandate, average defender posture quickly rises above 0.70. Although not any higher than the peak in Fig. 9a, the high posture is achieved much earlier in gameplay (timestep $t = 10$ instead of $t = 600$). The result is that defenders are able to achieve high posture while retaining high net worth early in simulation, which makes them very expensive to attack. Attackers—who have not yet been able to grow their wealth—do not have enough wealth to attempt attacks, and % attacked and % ransomed quickly drop to 0. With many rounds of no attacks, the game is considered to be at equilibrium and ends shortly after 50 rounds.

Another interesting finding is that a 1% resource mandate does not seem to remove the incentive for additional security investments; we also see that at $M = 1\%$ induces more demand for insurance as well (Fig. 10a). By the time $M = 2\%$, defenders' posture is likely high enough to allow them to buy neither insurance nor additional security.

Finally, we simulated higher values of M but omit the results here because the trends closely mimic the $M = 2\%$ case. Namely, defenders' posture quickly reaches $\bar{\mu_p} > 0.70$; the number of attacks per round quickly approaches 0 and the game quickly settles at an equilibrium; defenders retain a modest fraction of initial wealth, but begin to spend more and more on security (unnecessarily) and overall loss increases.

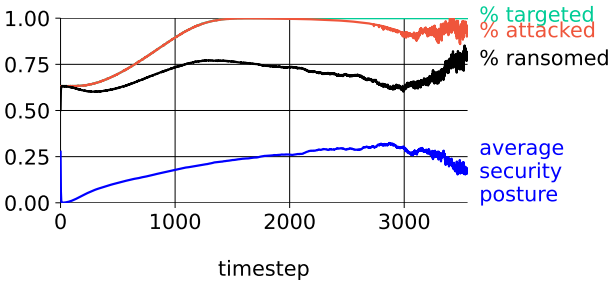

Figure 11: Gameplay barometer values given an insurance mandate. Compared to the baseline model (see Fig. 5), average defender security posture is lower.

Discussion: We highlight two takeaways from this study: first, we find that losses are minimized when defenders are required to invest at least 2% of resources into security. We note that this is roughly double the current real-world standard practice (as found in §3.6).

Perhaps the most significant finding from this study is that we find evidence that security is a weakest-link game *despite not including this as a feature of our model*[4]. In other words, it is an *emergent property* of our model that the existence of weak defenders determines the outcomes for strong defenders too. The mechanism by which this is possible is in the attackers' emergent strategy to attack weak defenders first and bootstrap themselves into being able to eventually take down all defenders.

[4]In game theory, a weakest-link (or minimum effort) game is cooperative game wherein the collective outcome for a group is determined by the player who puts in the least amount of effort towards some shared goal.

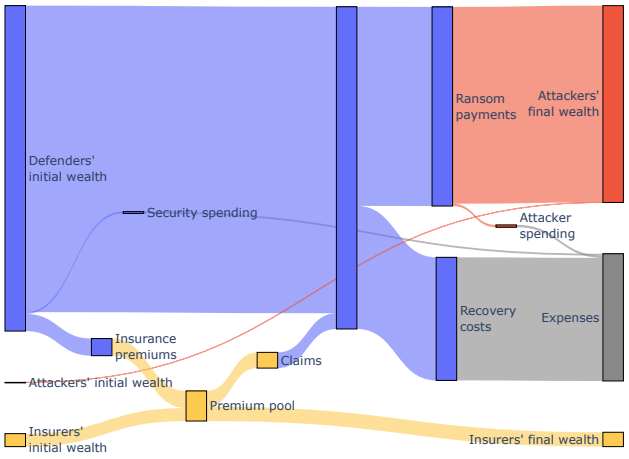

Figure 12: Wealth transfers given an insurance mandate. Compared to the baseline model (see Fig. 6), insurers' wealth increases by roughly 10%, but for the defenders, the outcome remains a total loss.

8 Mandated Insurance

Conventional insurance can be beset by the adverse selection problem, where parties with a high risk of loss are more likely than low-risk parties to seek insurance, causing risk pools to contain disproportionately risky policyholders [50]. In the presence of information asymmetry, insurers may not be able to distinguish high-risk policyholders from low-risk policyholders and may be forced to raise premiums for all policyholders in response. This can raise the price of insurance for all parties (including low-risk policyholders) and distort the insurance marketplace and lead to market failure. One potential solution to the adverse selection problem has been an insurance mandate requiring all parties to be insured [51].

There may be reason to believe that adverse selection besets cyber insurance as well: for example, insurers appear to have markedly different levels of sophistication when it comes to estimating insurees' risk levels [30]; this might allow for an information asymmetry between insurer and insuree and cause adverse selection. In our baseline model, such adverse selection is possible because defenders and insurers have different levels of knowledge on defenders' security posture (see §4.3.3).

Hence in our next study, we implement an insurance mandate policy requiring all defenders to obtain insurance. As with health care, the argument might be that requiring universal coverage increases the average cyber posture of the risk pool, allowing for a more functional insurance market.

Study Configuration: At the start of each round, each defender is required to purchase an insurance policy. The defender is able to choose the best quote of $Q = 10$ insurers, chosen at random. After purchasing the policy, defenders may choose to also make security investments.

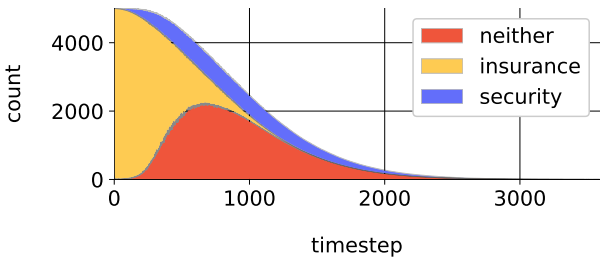

Figure 13: Defender choices when the insurance loss ratio $LR = 100\%$, indicating actuarially fair insurance.

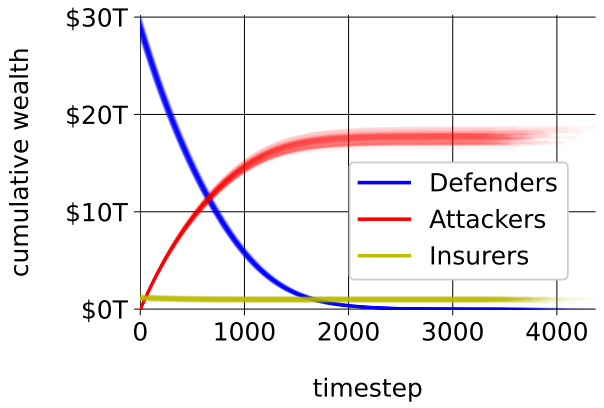

Figure 14: Cumulative assets under actuarially fair insurance. Insurers' wealth decreases by about 10% despite selling more policies.

Study Results: Somewhat surprisingly, we observe that under the existence of an insurance mandate, game outcomes do not strongly differ from the baseline model (Figs. 11, 12). The only notable difference was that insurers ended simulations with roughly 10% more than when they started (compared to the baseline model where the insurers actually lost wealth as a result of underwriting policies).

9 Actuarially Fair Insurance

In the baseline model we observe that there is not a strong incentive for defenders to purchase insurance (see Fig. 4). One explanation might be that premiums are priced too high relative to what claimants expect to receive in return. To test this hypothesis, we model and simulate an actuarially fair insurance market where insurers write policies such that expected gains are equal to expected losses. This may be analogous to various U.S.-based government-backed insurance schemes such as the Federal Crop Insurance Corporation (FCIC) or the National Flood Insurance Program (NFIP) [52, 53].

Study Configuration: We can model actuarially fair insur-

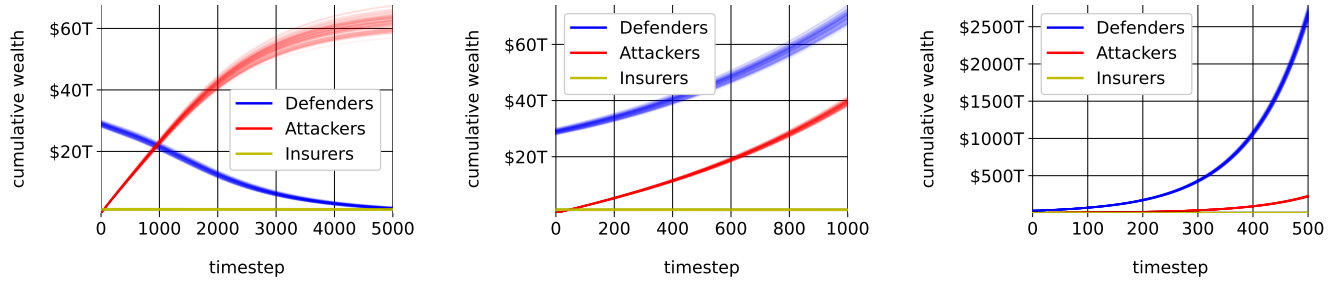

(a) $r = 0.001$ (b) $r = 0.002$ (c) $r = 0.01$

Figure 15: Cumulative agent wealths for various fixed growth rates. Each subfigure plots 100 simulations.

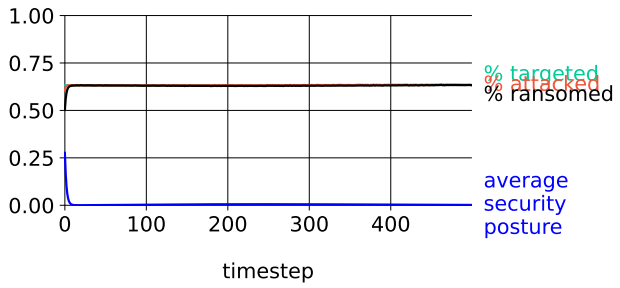

Figure 16: Gameplay barometer values when the growth rate $r = 0.01$. For all $r \geq 0.01$ the system reaches a homeostasis where posture drops to 0 and % targeted/attacked/ransomed stabilize at 0.63.

ance in our model by fixing the loss ratio to be LR=100%. This implies that insurers write policies with the expectation that the amount they gain from premiums is equal to the amount they will lose to claims.

Study Results: When *LR* increases to 100%, the expected loss given insurance decreases (see Eq. 39) making insurance a more attractive option (Fig. 13). Surprisingly, even given the increase in premiums collected, the insurers actually end up with *less* cumulative wealth than in the default game scenario (Fig. 14 vs. Fig. 3); This is because there is no profit buffer to insulate insurers from the consequences of poor underwriting when collecting defender security posture. We observe that despite the considerable change in defender behavior, model outcomes remains nearly unchanged from the baseline model evaluation.

It is not altogether unexpected that selfless insurers did not improve outcomes for defenders: increasing the incentive to purchase insurance does not affect the incentive to invest in security but instead replaced the incentive to do neither (Fig. 14 vs. Fig. 3). With defender posture mirroring the default configuration, there is no change in incentives for the attackers and no change in outcomes either.

10 A Non-Closed Ecosystem

An obvious limitation of our model so far is that it is a closed ecosystem where wealth is only transferred and never created. This is a clear deviation from the real world, where organizations earn profits and existing wealth experiences compound growth. In an effort to more accurately model the defender-attacker-insurer ecosystem, one may wish to add compound growth to our model.

We hesitate to include such growth into our model for two reasons: First, our interests thus far have been in measuring model outcomes, which requires a notion of model equilibrium or homeostasis (see §4); continuous wealth growth runs counter to this goal. Second, to keep the model as interpretable as possible, we wish to constrain our model to be a minimum viable representation of the security attack and investment ecosystem. Our concern is that additional model inputs—while in some sense making the model more "accurate"—may instead obfuscate model behavior, particularly defender decision-making. Nevertheless, in our effort to make our model a useful tool for others, we have included such a parameter: the growth rate r. This study explores the effects of compound growth.

Study Configuration: We introduce a new game parameter r which is the rate of return on defenders' assets. At the start of each round, each defender's wealth grows by $d_{i,w}^{[t]}(1+r)$.

Study Results: For small values of r, defenders are still fully looted but with greater wealth transferred to the attackers (Fig. 15a). As r increases, defenders are able to grow at roughly the same rate that attackers are able to ransom it away (Fig. 15b). For values of $r \geq 0.01$, defenders' growth is unburdened by attackers' ransoms. This is in spite of the trend for $r \geq 0.01$, where posture drops to 0 (Fig. 16).

We find that the model is highly sensitive to the growth rate and that the growth rate can dominate the system's behavior. Out of caution we decline to draw other strong conclusions from this study, but present it to demonstrate the flexibility of our model and its implementation.

11 Other Limitations

As noted above in §10, most of our studies were conducted within a closed ecosystem where wealth was only transferred and not created (although this limitation is self-imposed). In §3 we also mention our decision to model all security investment through a single investment payoff function rather than modeling specific security controls, which may be seen as a limitation or as a necessary abstraction. Other such tradeoffs were made elsewhere in our model as well: First, there is no interaction or interdependence between defenders or attackers. This makes our model unable to explicitly examine network effects between players or model how a security failure by one organization can compromise security in another organization[5]. We also do not model other insurance features like exclusions, third-party liability, or systemic risk in an effort to constrain model complexity.

Finally, we point out that much of the empirical work done in this paper was best-effort, and that all estimations and regressions may be improved by finding and including additional data sources.

12 Conclusion

In a world where evaluating security policies ex ante is not generally possible, our work serves as a blueprint for what may be the next best option: large-scale simulation of the security ecosystem using a game theoretic agent-based model.

Using our model, we are able to simulate various policy "interventions", first via a sensitivity analysis, and then as specific interventions like mandatory minimum investments, mandatory insurance, and actuarially fair insurance. From these studies we find that even perfect play from defenders fails to produce a social optimum.

This carries significant implications: security cannot be solved by expecting defenders to be better at defense. We find evidence that the lack of security standards creates easy money for attackers; over time this compounds until even the well-defended are victimized. Indeed, our model finds that security is a weakest-link game. This helps make the case for policy interventions in security, whereby setting minimum standards can disincentivize attackers and reduce losses. To make our work accessible to others, we have released it as an online application and have made the code open source.

Acknowledgments

We would like to thank the anonymous peer reviewers and shepherd for their feedback. We would also like to thank Sasha Romanosky, Daniel Woods, and Jason Healey who provided feedback on this work. Simha Sethumadhavan has a significant financial interest in Chip Scan Inc.

[5]Regardless, our defenders are imbued with the ability to react to the overall environment (Eq. 27) which does permit a level of interdependence as seen in §7.

Ethics Considerations

We did not encounter any ethical concerns while conducting this work. We do not believe this work has the potential to have any negative potential outcomes on any stakeholders.

Open Science Policy

We have released artifacts of this work in three formats: First, we provide a stable persistent release of artifacts via Zenodo, where we include all datasets, code, and figures used in this work [15]. Second, we provide a link to a GitHub repository which contains the same codebase as the Zenodo artifacts but may continue to be updated past the date of the Zenodo upload [16]. Third, as mentioned previously, we provide a link to a web hosted version of the simulator where users can easily interact with the model and adjust model inputs [11].

References

[1] "National cybersecurity strategy," The White House, 2023.

[2] "New smart devices cyber security laws one step closer," Department for Digital, Culture, Media and Sport, United Kingdom, 2022.

[3] "Proposal for a regulation of the european parliament and of the council on horizontal cybersecurity requirements for products with digital elements and amending regulation (EU) 2019/1020," European Commision, European Union, 2022.

[4] "Cybersecurity certification guide," Cybersecurity Certification Centre, Cyber Security Agency of Singapore, 2021.

[5] "Statement of compliance for the cybersecurity label," National Cyber Security Centre, TRAFICOM, Finland, 2019.

[6] R. Anderson, "Why information security is hard-an economic perspective," in *Seventeenth Annual Computer Security Applications Conference*, pp. 358–365, IEEE, 2001.

[7] R. Anderson and T. Moore, "The economics of information security," *Science*, vol. 314, no. 5799, pp. 610–613, 2006.

[8] A. Hastings and S. Sethumadhavan, "WaC: A new doctrine for hardware security," in *Proceedings of the 4th*

ACM Workshop on Attacks and Solutions in Hardware Security, pp. 127–136, 2020.

[9] H. W. Rittel and M. M. Webber, "Dilemmas in a general theory of planning," *Policy sciences*, vol. 4, no. 2, pp. 155–169, 1973.

[10] D. Clemente, "International security: Cyber security as a wicked problem," *The World Today*, vol. 67, no. 10, pp. 15–17, 2011.

[11] https://cyberspending.cs.columbia.edu/

[12] D. W. Woods and T. Moore, "Does insurance have a future in governing cybersecurity?," *IEEE Security & Privacy*, vol. 18, no. 1, pp. 21–27, 2019.

[13] J. MacColl, J. R. Nurse, and J. Sullivan, "Cyber insurance and the cyber security challenge," *RUSI Occasional Paper*, 2021.

[14] J. Wolff, *Cyberinsurance policy: Rethinking risk in an Age of ransomware, computer fraud, data breaches, and cyberattacks*. MIT Press, 2022.

[15] A. Hastings, "Artifacts for usenix security symposium '25 paper "voluntary investments, mandatory minimums, or cyber insurance: What minimizes losses?"."

[16] https://github.com/columbia-castl/monte-carlo-security-games/

[17] L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Transactions on Information and System Security (TISSEC)*, vol. 5, no. 4, pp. 438–457, 2002.

[18] R. Böhme and G. Kataria, "On the limits of cyber-insurance," in *International Conference on Trust, Privacy and Security in Digital Business*, pp. 31–40, Springer, 2006.

[19] J. Grossklags, N. Christin, and J. Chuang, "Secure or insure?: A game-theoretic analysis of information security games," in *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, (New York, NY, USA), pp. 209–218, ACM, 2008.

[20] R. Böhme, G. Schwartz, *et al.*, "Modeling cyber-insurance: towards a unifying framework.," in *WEIS*, 2010.

[21] R. Böhme, T. Moore, *et al.*, "The "iterated weakest link" model of adaptive security investment," *Journal of Information Security*, vol. 7, no. 02, p. 81, 2016.

[22] R. Pal and L. Golubchik, "Analyzing self-defense investments in internet security under cyber-insurance coverage," in *2010 IEEE 30th international conference on distributed computing systems*, pp. 339–347, IEEE, 2010.

[23] R. Pal, L. Golubchik, K. Psounis, and P. Hui, "Will cyber-insurance improve network security? A market analysis," in *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*, pp. 235–243, IEEE, 2014.

[24] R. Pal, L. Golubchik, K. Psounis, and P. Hui, "Improving cyber-security via profitable insurance markets," *ACM SIGMETRICS Performance Evaluation Review*, vol. 45, no. 4, pp. 7–15, 2018.

[25] T. Yin, A. Sarabi, and M. Liu, "Deterrence, backup, or insurance: game-theoretic modeling of ransomware," *Games*, vol. 14, no. 2, p. 20, 2023.

[26] T. Meurs, E. Cartwright, and A. Cartwright, "Double-sided information asymmetry in double extortion ransomware," in *International Conference on Decision and Game Theory for Security*, pp. 311–328, Springer, 2023.

[27] S. Erol and M. J. Lee, "Financial system architecture and technological vulnerability," in *Staff Reports no. 1122*, Federal Reserve Bank of New York, October 2024.

[28] A. Baldwin, I. Gheyas, C. Ioannidis, D. Pym, and J. Williams, "Contagion in cyber security attacks," *Journal of the Operational Research Society*, vol. 68, no. 7, pp. 780–791, 2017.

[29] X. Xie, C. Lee, and M. Eling, "Cyber insurance offering and performance: An analysis of the us cyber insurance market," *The Geneva Papers on Risk and Insurance-Issues and Practice*, vol. 45, pp. 690–736, 2020.

[30] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones, "Content analysis of cyber insurance policies: How do carriers price cyber risk?," *Journal of Cybersecurity*, vol. 5, no. 1, p. tyz002, 2019.

[31] D. W. Woods, T. Moore, and A. C. Simpson, "The county fair cyber loss distribution: Drawing inferences from insurance prices," *Digital Threats: Research and Practice*, vol. 2, no. 2, pp. 1–21, 2021.

[32] I. Kotenko, "Agent-based modeling and simulation of cyber-warfare between malefactors and security agents in internet," in *19th European Simulation Multiconference "Simulation in wider Europe*, 2005.

[33] S. Hofmeyr, T. Moore, S. Forrest, B. Edwards, and G. Stelle, "Modeling internet-scale policies for cleaning up malware," in *Economics of Information Security and Privacy III*, pp. 149–170, Springer, 2013.

[34] D. Woods and A. C. Simpson, "Monte carlo methods to investigate how aggregated cyber insurance claims data impacts security investments," 2018.

[35] S. Matsugaya, "Rise in active raas groups parallel growing victim counts," Trend Micro, March 2024.

[36] "The state of ransomware 2024," Sophos, 2024.

[37] "Report on the cyber insurance market," National Association of Insurance Commissioners, 2022.

[38] companiesmarketcap.com, 2024.

[39] I. Ilascu, "REvil ransomware gang claims over $100 million profit in a year," BleepingComputer, 2020.

[40] L. Constantin, "REvil ransomware explained: A widespread extortion operation," 2021.

[41] "U.S. and U.K. disrupt lockbit ransomware group and indict two russian nationals while OFAC levies sanctions," Chainalysis, 2024.

[42] I. W. Gray, J. Cable, B. Brown, V. Cuiujuclu, and D. McCoy, "Money over morals: A business analysis of Conti ransomware," in *2022 APWG Symposium on Electronic Crime Research (eCrime)*, pp. 1–12, IEEE, 2022.

[43] T. Puech and L. Laurenne-Sya, "Ransomware: Inside the former CONTI group," RiskInsight, 2022.

[44] "State of the tech spend pulse," Flexera, 2022.

[45] "Report shows cybersecurity budgets increased 6% for 2022-2023 cycle," Security Magazine, Sept. 2023.

[46] "Cybersecurity insights 2023: Budgets and benchmarks for financial services institutions," Deloitte, 2023.

[47] B. F. Ginos, *Parameter estimation for the lognormal distribution*. Brigham Young University, 2009.

[48] R. P. Brent, "An algorithm with guaranteed convergence for finding a zero of a function," *The computer journal*, vol. 14, no. 4, pp. 422–425, 1971.

[49] J. Easterly and E. Goldstein, "Stop passing the buck on cybersecurity," Foreign Affairs, 2023.

[50] M. Rothschild and J. Stiglitz, "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information," in *Uncertainty in economics*, pp. 257–280, Elsevier, 1978.

[51] M. B. Hackmann, J. T. Kolstad, and A. E. Kowalski, "Adverse selection and an individual mandate: When theory meets practice," *American Economic Review*, vol. 105, no. 3, pp. 1030–1066, 2015.

[52] "Federal crop insurance: A primer," Congressional Research Service, Feb. 2021.

[53] E. O. Michel-Kerjan, "Catastrophe economics: the national flood insurance program," *Journal of economic perspectives*, vol. 24, no. 4, pp. 165–186, 2010.