Cyberinsurance and Public Policy: Self-Protection and Insurance with Endogenous Adversaries

Fabio Massacci[a,*], Joe Swierzbinski[c], Julian Williams[b]

[a] *Department of Information Engineering and Computer Science, University of Trento, Italy.*
[b] *Durham University Business School, Durham, UK.*
[c] *University of Aberdeen Business School, Aberdeen, UK.*

Abstract

Corporate insurance contracts providing liability coverage in the event of an information security breach are increasingly popular. In addition to the obvious use of 'Cyberinsurance' as a risk mitigation tool, a public policy narrative has emerged whereby insurance companies act as a clearing house for information and then provide guidance on appropriate security investment to firms seeking liability coverage. Utilizing few assumptions, our modeling framework demonstrates that this view of cyberinsurance as a delegated policy tool is unlikely to yield the anticipated coordination benefits, and may in fact erode the aggregate level of security investment undertaken by targets.

Keywords: Insurance, Cyber security, Public Economics, Optimal Investment Allocations
JEL Code: G1, L12, L22, L41, M21.

The market for corporate insurance contracts that provide liability coverage in the event of a security breach to a firm's IT systems ('cyber insurance') has been growing steadily since the turn of the millennium. As of 2015, this market has grown to around $2.5B in annual premiums and estimates suggest that is could reach $7.5B before 2020. See O'Hearn and *et al.* (2015) for an overview.

From the perspective of a policy maker, both in the U.S. and in Europe, the anticipated benefits from a cyber insurance market are much higher than a simple financial instrument for controlling individual corporate risk exposures. For example, the chair of the U.S. Senate's Subcommittee on Consumer Protection, Product Safety, Insurance, and Data Security observed in March 2015:

> *While an insurer's primary function is to mitigate financial losses – not defend against cyber threats – cyber insurance may be a market-led approach to help businesses improve their cybersecurity posture by tying policy eligibility or lower premiums to better cybersecurity practices.*

*Corresponding Author. The authors would like to thank Luca Allodi from the University of Trento, Vadim Kotov from Bromium, and the members of the Computer Laboratory in Cambrige (in particular Ross Anderson, Richard Clayton, Daniel Thomas, and Sultan Kus) for very useful discussions and insights on hackers' technology and markets. We would like also to thank the participants to the Lorentz' Adversarial Risk Analysis seminar (in particular Milind Tambe, Wolter Pieters, Vivian Jacobs, David Banks, Dieter Gollmann, Andr Hoogstrate, and Christian Probst) for useful discussions on the use of game theory techniques for security, Angela Sasse and her group at UCL, Alex Ashby from Oxford, Christos Ioannidis from the University of Bath, and the seminar participants at the University of Durham (in particular Parantap Basu, Abderrahim Taamouti, Hugo Kruiniger, Leslie Reinhorn, Xiaogang Che, and Damian Damianov) for useful comments. Any remaining mistakes are the sole responsibilities of the authors.

Email addresses: `fabio.massacci@unitn.it` (Fabio Massacci), `j.swierzbinski@abdn.ac.uk` (Joe Swierzbinski), `julian.williams@durham.ac.uk` (Julian Williams)

Source: U.S. Senate 2015.[1]

furthermore

> *A robust cybersecurity insurance market could help reduce the number of successful cyber attacks by: (1) promoting the adoption of preventative measures in return for more coverage; and (2) encouraging the implementation of best practices by basing premiums on an insureds level of self-protection.*

Source: U.S. Department of Homeland Security 2017.[2]

Similar expectations are held by several high profile cyber security experts such as (Schneier, 2001, Chapter 5) or Gordon et al. (2003), as well as leading economists (Varian, 2000).

A natural model of multiple target firms, active cyber adversaries, policy maker, and insurance companies presented in this paper demonstrate that the presence of attackers reacting to groups of firms security investments generates non-linear dependencies between firms. More strikingly, it shows that the presence of insurance contracts, either actuarially fairly or unfairly priced, result's in a reduction of aggregate security investment and, hence, an increase in the number of active attackers. Our results contribute to the growing literature on how to describe interdependencies between agents and the impact on public policy. See Bramoullé et al. (2014) and Allouch (2015) for recent examples using networks and linear payoff functions. In our case, the dependency–mechanism, as a result the joint action of attackers and targets, exhibits substantial non-linearities that are hard to capture with a standard network model with linear pay-offs.

From a classical perspective, firms with well diversified owners should not buy liability insurance as managers should reflect the risk neutrality of the diversified owners (Mayers and Smith Jr, 1987). However, the large corporate liability insurance market and the heavy demand for financial hedging instruments indicate that firm decision making is most likely made under a risk averse basis, driven by the corporate officers and their incentives, as argued by Grossman and Hart (1982). Indeed, Caillaud et al. (2000), propose a model whereby that risk-neutral firms are 'induced to risk-aversion' because the disclosure of 'accidental-losses' that deteriorate the profitability of projects is private to the firm and costly audit is needed to demonstrate that the 'accident' was something that the agents managing the firm could not control through an appropriate investment, under a risk neutral marginal cost/marginal benefit trade-off. Furthermore, MacMinn and Garven (2000) argue more directly that the demand for corporate insurance stems from the firm choices mimicking the risk averse behavior of its corporate officers, who have significant costs attached to adverse events through imperfections in the executive labor market and other reputational effects.

The notion corporate officers buy insurance to hedge against risks to their own positions is empirically investigated from a legal perspective in Baker and Griffith (2007), who use the demand for corporate liability insurance as a predictor of firms bad corporate governance. Indeed, Griffith (2006) argues that SEC should mandate insurance details of officers and

[1] See the minutes of the hearings at `http://www.commerce.senate.gov/public/index.cfm/2015/3/examining-the-evolving-cyberinsurance-marketplace.`, in addition the Network and Information Security (NIS) Directive of the European Union, looks at the role of Cyber insurance in providing forward guidance see `http://www.cspforum.eu/uploads/1_Cyber-insurance_and_NISD.pdf` and `https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive`.

[2] The full quote is from: `https://www.dhs.gov/cybersecurity-insurance`, last accessed February 20, 2017. Paragraph 1:2.

directors liability insurance policies in order to allow investors to police corporate management. Hence, there is a substantial case that, in the absence of insurance, risk management decisions are made on a risk-averse basis, hence the observed demand for liability insurance.

The impact of moral hazard and adverse selection in the presence of insurance has a long history of investigation in economics – see for instance (Pauly, 1974; Shavell, 1987; Cornes and Sandler, 1996; Freeman and Kunreuther, 1997) for an eclectic set of examples that directly relate to the notions of public policy, liability sharing, and insurance considered herein.

How the presence of insurance and the collective behavior of victims of crimes can influence the aggregate behavior of the criminals who generate the risks being insured has been less comprehensively investigated in the literature. Dionne and Wang (2013) have shown that external environmental conditions affect the distribution of insurance claims and their legitimacy in the case of auto insurance fraud. However, Dionne and Wang (2013) suggest that the systematic factor that varies across all targets is driven by an external macroeconomic effect and is not fully endogenous.

Some discussion has occurred in the insurance literature on theft – for instance if large numbers of households buy burglar alarms, the aggregate cost of being a burglar increases as one needs to invest in specialist skills to bypass alarms or to spend more time searching for vulnerable homes (Kunreuther and Heal, 2003b). Fewer burglars should then join the market for burglaries and the actuarial risk of a payout decreases. However, the job selection choice for burglars is likely to be highly inelastic; there are high costs associated with changing career and the stigma of prior convictions can result in substantial costs when choosing an alternative career path. Criminologists commonly refer to the concept of 'consistency' in behavior. That is, once a pattern of offending behavior has been established it is very difficult for the offender to adjust to changing opportunity sets and costs associated with this behavior.

Kirwan and Power (2013) indicate that individuals and organizations engaged in cyber crime do not face such difficulties. For example, the choice of a security engineer to work on either malicious software or software with a legitimate business purpose is simply a matter of re-tasking oneself.[3] See also the discussion in Miller (2007); Johnson (2014) and McCarthy (2002).

In this paper, we outline a game played between a group of targets who invest in defensive expenditure to reduce the risk of a successful attack by one or more attackers. Attackers are modeled as agents in competition with upfront entry costs who infiltrate and expropriate valuable information from the targets. Into this game we introduce a variety of insurance contracts and policy driven behavioral restrictions that the insurer can impose on the target. Our main result is quite surprising. When losses are significant neither the monopolist insurers nor a fully competitive insurance market have the incentive to reduce externalities within the market. Furthermore, a rational monopolist insurer would be positively assuaged towards inflating the cyber threat by mandating insufficient security expenditures, as long as it can identify the actuarially fair price of insurance risk and the maximum quote premium it can charge to risk averse targets. The fact that attackers react to aggregate investment changes the opportunity set for a monopolist, hence the effect is almost exclusive to situations when attackers react relatively quickly to adjustments in target investment.

The next section, §(1), outlines the mathematical description of target firms, attackers,

[3]The recent disclosure of the emails of the Italian "HackingTeam" organization showed that the same software and services were sold to the FBI, the Italian Police, Russian Hackers, and the Sudanese Intelligence. See `http://www.theguardian.com/technology/2015/jul/06/hacking-team-hacked-firm-sold-spying-tools-to-repressive-regimes-documents-claim`.

insurers, and a public policy coordinator. In §(2), we present baseline equilibria where heterogeneous targets self protect with and without a policy coordinator. §(3) compares the self-protection equilibria in several insurance market models including a monopoly insurer and actuarially fair insurance with a public policy coordinator. We illustrate our general results with two numerical examples in §(4) In one example, a continuum of identical firms is used to explain the insurance trap graphically. In the other example, a continuum of firms of two types (large and small) is used to show that the potential response of attackers to the perceived vulnerability of targets can cause large changes in security expenditures. We provide a general summary and some natural extensions in §(**?**).

1. Modeling Assumptions

We use fully-informed agents for two reasons. First, there is no consensus in the computing science literature on who are the better informed parties and what are the types of information asymmetries that occur. Second, we are interested in exploring the new public policy issues raised by our model in the simplest plausible setting before adding complications.[4]

The timeline of the games is as follows: for our first game, in the absence of insurance or a regulator, targets and attackers simultaneously choose their optimal expenditure and participation in the "market for attacks". In game two, a public policy coordinator initially sets a mandatory investment profile and in the second stage attackers choose their optimal participation.

In game three, targets simultaneously choose whether or not to buy an actuarially fair insurance contract as well as their levels of security expenditure. At the same time, attackers choose whether or not to participate in an attack given their forecast of the profile of security expenditures. We make the standard assumption that a competitive insurance market with full information will result in the provision of actuarially fair insurance policies without further modelling the behaviour of insurers.

In game four, we combine games two and three to illustrate the joint impact of regulation and insurance. As in game two, the public policy coordinator mandates the security investment for each target in the first stage of the game. In the second stage, targets simultaneously decide whether or not to purchase actuarially fair insurance in a competitive market while attackers simultaneously decide whether or not to participate in an attack.

Finally, in the first stage of game five, we relax the actuarially fair insurance assumption and model a monopoly insurer making a take-it-or-leave-it offer of an insurance contract tailored for each target. In the second stage of the game, targets simultaneously decide whether or not to accept the monopolists offer. If targets accept the contract then their level of security expenditure is set by the contract. If a target rejects the monopolists offer, then the target also chooses its level of security expenditure. Also in the second stage, attackers simultaneously choose whether or not to participate in an attack.

1.1. Targets Assumptions

Each target $i \in \{1 \dots N\}$ has an endowment of assets $W_i > 0$ and, in the event of a successful cyber attack, is subject to losses L_i where $L_i > 0$. Each target makes a security investment x_i, where $0 \leq x_i < L_i$ and is targeted by a number of attackers n_i where

[4]Without irony Pal et al. (2013) provides a summary of how a cyber security vendor acting both as security provider and as insurer can extract a sizeable rent, even under symmetric information on the actual threat level.

$0 \leq n_i < \infty$. The probability of a successful attack against target i is determined by its investment x_i and the number of attackers n_i focusing on this particular target, and is denoted by $\sigma_i(n_i, x_i)$. For each target there are only two outcome states, "successfully-attacked" and "not-successfully-attacked"; the probability of each state is respectively σ_i and $1 - \sigma_i$.

From this point onwards, we follow Grossman and Hart (1982) and Caillaud et al. (2000) and when we refer to "target-preferences" or "target-utility" we are referring to the induced preferences exhibited by the firm reflecting the "transferred" preferences from the corporate officers to the revealed actions of the firm. We are, of course, not implying that firms have behavioral characteristics as individuals do.

[$\mathcal{A}$.1] The utility function of the $i \in \{1, \ldots, N\}$ firm generated by the preferences of their corporate officers is denoted $U_i(w)$ over a scalar random wealth variable $w \geq 0$. $U_i(w)$ is an least twice differentiable von Neumann Morgenstern utility function, such that for all w, $U_i'(w) > 0$ and $U_i''(w) < 0$ which are respectively the first and the second derivatives of $U_i(\cdot)$ w.r.t. w.[5]

For compactness we define the following notation convention, let $\mathscr{F}(.)$ be an at least twice differentiable continuous function with vector of arguments $\boldsymbol{\theta} = [\theta_j]$ we utilize extensively an elasticity operator, $\mathrm{e}_{\mathscr{F}(\boldsymbol{\theta}),\theta_i} := (\mathscr{F}(\boldsymbol{\theta})/\theta_i)/(\partial\mathscr{F}(\boldsymbol{\theta})/\partial\theta_i)$. Furthermore, when θ_i is not explicitly defined we presume that the elasticity is in terms of x_i. Hence, $\mathrm{e}_{\Delta U_i(x_i)} = \mathrm{e}_{\Delta U_i(x_i),x_i}$. Finally, in certain cases we need to define a second order elasticity of the form $\mathrm{e}_{\mathscr{F}'_{\theta j}(\boldsymbol{\theta}),\theta_i} := (\mathscr{F}'_{\theta j}(\boldsymbol{\theta})/\theta_i)/(\partial\mathscr{F}'_{\theta j}(\boldsymbol{\theta})/\partial\theta_i)$. Where $\mathscr{F}'_{\theta j}(\boldsymbol{\theta}) = \partial\mathscr{F}(\boldsymbol{\theta})/\partial\theta_j$. Again, for compactness of exposition, if θ_i is not explicitly defined, for instance $\mathrm{e}_{\mathscr{F}'(\boldsymbol{\theta})}$ or then it is presumed that $\theta_i = \theta_j = x_i$, hence a second order elasticity in x_i.

The expression $\Delta U_i(x_i) = U_i(W_i - x_i) - U_i(W_i - x_i - L_i)$ denotes the drop in utility between the two outcome states for a given investment x_i. Notice that $\sigma_i \Delta U_i$ captures the expected drop in utility in case of a successful attack. The elasticity $\mathrm{e}_{\Delta U_i}$ with respect to the security investment x_i intuitively captures the change in utility for an upfront investment, a sure cost x_i from the perspective of a target, to face a possible loss L_i. For risk averse targets ($\mathcal{A}$.1), it is $\mathrm{e}_{\Delta U_i} > 0$.

Given wealth w_i and an expected loss $\sigma_i L_i$, we follow (Pratt, 1964, p.124) and define the risk premium π_i, see (Pratt, 1964, p.124), as the amount that would make a risk averse target indifferent between incurring the certain loss $\pi_i + \sigma_i L_i$ and being exposed to the uncertain loss L_i with probability σ_i. We can define the maximum risk premium $\pi_{\max}(w_i, L_i)$ as the maximum of the risk premium π_i over all σ_i. The corresponding probability of loss where such maximum premium is attained is denoted by $\sigma_{\max}(w, L_i)$, and the maximally insurable loss $\mathfrak{L}(w_i, L_i)$ is the sum of the two:[6]

$$U_i(w_i - \pi_{\max} - \sigma_{\max} \cdot L_i) = \sigma_{\max} \cdot U_i(w_i - L_i) + (1 - \sigma_{\max}) \cdot U_i(w_i) \tag{1}$$

$$\mathfrak{L}(w_i, L_i) = \pi_{\max}(w_i, L_i) + \sigma_{\max}(w_i, L_i) \cdot L_i \tag{2}$$

For a risk averse target, under our assumptions, the existence of $\sigma_{\max}(w_i)$ and $\pi_{\max}(w_i)$ is guaranteed by Jensen's inequality and the mean value theorem.

For a risk neutral target, the risk premium pi_i is equal to zero, for all values of w_i, σ_i and L_i. Hence, $\pi_{\max}(w_i) = 0$ and $\sigma_{\max}(w_i)$ can take any value in $[0, 1]$. Note that, in either case,

[5] Corporate liability insurance is expected to cover large losses (O'Hearn and *et al.*, 2015) and therefore the critique of expected utility theory by Rabin (2000) may not apply as it concerns decisions for lotteries with small stakes.

[6] (Pratt, 1964, p.124) refers to the sum $\pi_i + \sigma_i L_i$, without maximization, as the 'insurance premium'.

the maximum risk premium, the quantity $\sigma_{\max}(w_i)$ and hence, the the maximally insurable loss do not depend on the actual probability of loss σ_i.

For an individual firm, there may only be a discrete set of possible investments in security controls. For the population of targets, the assumption of a continuous probability function ensures tractability and ease of exposition. Hence, we assume that the probability of a successful attack $\sigma_i(x_i, n_i)$ against target i is at least twice differentiable in x_i and n_i.

The properties described below are well understood properties for implementing security controls. Once simple controls are implemented, the cost for marginal improvement in security will increase, with increasing rapidity (Gordon and Loeb, 2002).

[$\boldsymbol{\mathcal{A}}$**.2**] The quantity $\sigma_i(\cdot)$ is a strictly increasing function of the number of attackers; that is $\forall n_i$ $(\partial\sigma_i/\partial n_i > 0)$. In the absence of attackers the probability of a successful attack is zero: $\forall x_i$, $n_i = 0$ implies $\sigma_i = 0$, .

[$\boldsymbol{\mathcal{A}}$**.3**] For all $n_i > 0$, the quantity $\sigma_i(\cdot)$ is continuous strictly decreasing with increasing investment by target i in security investment, for all $x_i \geq 0$ $(\partial\sigma_i/\partial x_i < 0)$. The marginal probability of a successful attack with respect to the number of attackers also decreases with increasing security expenditure $(\partial(\partial\sigma_i/\partial n_i)/\partial x_i \leq 0)$.

[$\boldsymbol{\mathcal{A}}$**.4**] The rate of reduction in $\sigma_i(\cdot)$ with increasing x_i is strictly decreasing with increasing defensive expenditure, for all $x_i \geq 0$ $(\partial^2\sigma_i/\partial x_i^2 > 0)$.

In our approach, externalities between targets are entirely driven by the aggregate reaction of attackers n_i. It is possible to extend the formal treatment by imposing an explicit externality which directly link the collection of investments $\mathbf{x}_{-i} = [x_j]_{j\neq i}$ of other targets to the probability σ_i of a successful attack on target i. This approach is pursued by Kunreuther and Heal (2003a,b) which explicitly considers externalities in the utility function in terms of all $\mathbf{x}$ whilst holding n constant. Our approach introduces a new channel by which externalities between targets can occur, and it seems sensible to first investigate this new effect in isolation without complicating matters further by adding direct dependencies between targets. Furthermore, our channel is a result of underlying behavior and does not rely on exogenous imposition of mutual dependencies between targets.

For the success probability, it is useful to consider two elasticities in the security expenditure x_i which capture the marginal effectiveness of security expenditures: the elasticity of the marginal increase in attack probability for increasing security expenditures $\mathrm{e}_{\partial\sigma_i/\partial x_i}$, and the elasticity of the marginal increase in attack probability for increasing number of attackers $\mathrm{e}_{\partial\sigma_i/\partial n_i}$. Both elasticities are less than zero as the probability diminishes with respect to increasing security expenditure. Their absolute value can be considered as an indication of the effectiveness of security expenditures. If the success probability can be decomposed into two separate multiplicative components $\sigma_i(x_i, n) = f_i(x_i) \cdot g_i(n)$, one depending only from n_i and one depending only from x_i, then the marginal elasticity in the increasing number of attackers coincides with elasticity of σ_i in x_i, that is $\mathrm{e}_{\partial\sigma_i/\partial n_i} = \mathrm{e}_{\sigma_i}$.[7]

1.2. Attacker Assumptions

Cyber attacks against each target i are attempted by a fraction of attackers n_i from a large pool of N_A potential attackers[8]. To mount an attack, an attacker must spend an

[7] A simple functional form that satisfies this condition is $\sigma_i = e^{-\alpha_i x_i} n_i^{\beta_i}$ where $\alpha_i > 0$ and $0 < \beta_i < 1$ are positive scalar parameters and $\sigma_i \in [0, 1]$ has an upper bound in $\ln n_i = \alpha_i/\beta_i x$. In this case we have $\mathrm{e}_{\sigma_i} = \mathrm{e}_{\partial\sigma_i/\partial x_i} = -\alpha_i$. Multiplicative forms have also been considered by Gordon and Loeb (2002); Kunreuther and Heal (2003a); Cavusoglu et al. (2008).

[8] Similarly to Ransbotham and Mitra (2009) we distinguish between 'advanced persistent threats' and 'large scale' attacks. The former are highly idiosyncratic risks are commonly covered under national security

upfront cost C_i and, if the attack is successful, realizes a reward R_i.[9]

Note that the attackers cost, C_i, and reward if successful, R_i, are allowed to depend on the target i but not on the identity of the particular attacker.

[$\mathcal{A}$.5] Attackers are risk neutral and make binary attack or no-attack decisions if the market conditions are favorable.[10] Each additional attacker against target i increases the saturation of the market and reduces the overall marginal benefit for all attackers: $\frac{\partial}{\partial n_i}\sum_{i=1}^{N}(R_i\ \sigma_i - C_i n_i) < 0$.

The second part of $\mathcal{A}$.5 does place a natural limit on the overall expected reward: attacks cannot be used as "money pumps" to extract arbitrarily large surpluses.

[$\mathcal{A}$.6] The attacker-target matching probability distribution is at maximum entropy.[11] Hence, the probability matching a given j attacker to the i target is random with a probability is $\varsigma_i = N_A/N$ and we can assume that $n_i = n$.[12]

[$\mathcal{A}$.7] In the event of a successful attack on target i the successful attacker does not share this reward with other attackers attacking the target and no further attack will generate any reward.[13] Further, when looking at the aggregate, the cost of the infrastructure can be assumed to be relatively constant $C_i = C$.[14] We denote the *rate of return* on C for a given reward R_i for an attack on the i target as $R_i/C = \rho_i$.

1.3. The Cyberinsurance Market

Let the i-th target have an available insurance contract described by the three-tuple (q_i, ℓ_i, x_i). The quantity q_i specifies the premium or quote paid upfront by i prior to any

---

rather than criminal or commercial liability governance. Further, recent studies estimate only few attacks are of the latter type Bilge and Dumitras (2012). Also in (Verizon, 2016, page 22), in spite of the front cover fanfare, cyber-espionage incidents totaled for a paltry 247 out of 64,199. In contrast, large-scale attacks affect individuals Grier et al. (2012), organizations Ransbotham and Mitra (2009), and industrial systems Nicholson et al. (2012) alike. Verizon (2016) reports that Crimeware and Web attacks alone represents over two thirds of the attacks of the financial sector, a figure aligned with the attack trend quantified by Google's researchers in Rajab et al. (2011).

[9] Attackers' rewards R_i are distinct from targets' losses L_i as they might be captured by different utility functions ranging from financial rewards (Hutchings and Clayton, 2016) to political gains (Li et al., 2011), from machines to be resold on black markets (Allodi et al., 2016) to kudos in forums (Ooi et al., 2012).

[10] The common knowledge of attack opportunities is due to the propensity of attackers to communicate success and failures through online forums Ooi et al. (2012); Allodi et al. (2016).

[11] Assumption $\mathcal{A}$.6 may appear restrictive from an economic perspective. However, exploits kits – representing two thirds of the threats against end users according to Google (Rajab et al., 2011) –, phishing attacks (Moore and Clayton, 2009), and most advanced persistent threats (Li et al., 2011) exploit the victim's clicks on a link in a web site, an email, or in a crafted document that is redirected to a place where malware is then provisioned to the visitor instead of 'normal' web content (Kotov and Massacci, 2013). Such technologies substantially reduce the costs of attack campaigns (Grier et al., 2012), but make it harder to predict who will eventually click on a link (ς_i) and be specifically targeted by the subsequent attack ($\sigma_i(\cdot)$). To address this uncertainty, the Rock Phish gang, a specialized 'enterprise' in phishing bank credentials, offered to its rogue 'clients' a server hosting fake websites from different banks (Moore and Clayton, 2007).

[12] One can think of each investment C_i as buying as a service a campaign of n_{C_i} attacks across targets (Grier et al., 2012) and the number of attackers per target n can be thought of as a product $n = N_A \cdot n_{C_i} N^{-1}$.

[13] Compromised goods (being them machines or credit cards) have very limited values by themselves on the black markets as the risk of buying 'scams for scammers' is high (Herley and Florêncio, 2010). A solid community reputation is needed by a seller to show that its machines are technically compromised but the reward has not been extracted yet (Allodi et al., 2016).

[14] Empirical studies demonstrate that exploit-kits only utilize a constant handful of exploits (Kotov and Massacci, 2013; Allodi and Massacci, 2014) and simple Google searches can be used to automated the reconnaissance of vulnerable websites to implant links (Moore and Clayton, 2009). For phishing attacks the cost of the harvesting kit is also limited and the largest cost is the spam campaign a bulk effort from a limited number of individuals (Moore and Clayton, 2008). In alternative, redirected traffic can be directly bought in bulk from malvertising providers as illustrated by Sood and Enbody (2011).

loss at the commencement of the contract. The quantity $\ell_i \leq L_i$ denotes the amount of the deductible or excess that will be left to be paid by target i if a successful attack against it occurs. Finally, x_i denotes the security standard that is mandated by the insurance contract as a minimum security expenditure required to qualify for insurance.[15]

[$\mathcal{A}$.8] Cyberinsurance companies are profit maximizers with a risk-neutral break-even requirement for the issuance of coverage at an actuarially fair price.

[$\mathcal{A}$.9] The security expenditures x_i are fully auditable, both ex ante and ex post, and the investment in defensive expenditure is made with commitment.[16]

[$\mathcal{A}$.10] The insurance company can identify the aggregate number of attackers per target for a given level of defensive expenditure across the population of targets, both in the presence or absence of insurance.[17] Therefore the provider of the insurance can fully determine the actuarially fair value of insurance for the i target in the presence or absence of insurance.

2. The Self-Protection Mechanism

2.1. Unregulated Markets

We start our treatment of the model with unregulated targets and and denote by $\mathbb{E}[U_i(x_i)|n]$ the expected utility of target i for a security investment x_i given an environment with n attackers per target:

$$\mathbb{E}[U_i(x_i)|n] = (1 - \sigma_i(x_i, n))\, U_i(W_i - x_i) + \sigma_i(x_i, n) U_i(W_i - x_i - L_i) \tag{3}$$

From $\mathcal{A}$.5 and $\mathcal{A}$.7 is follows that when n attackers attack a single target the first successful attacker extract the reward R_i and the probability of successfully attacking the i target is $n^{-1}\sigma_i(\cdot)$. At the equilibrium the expected profit of the n attackers per target aggregated over all possible targets must be equal to the aggregate costs of launching the attacks:

$$\sum\nolimits_{i=1}^{N} R_i \cdot \tfrac{1}{n} \cdot \sigma_i(x_i, n) \cdot \varsigma_i = \sum\nolimits_{i=1}^{N} c_i \cdot \varsigma_i \tag{4}$$

Whilst singularities may exist in the solution $n^\star(\mathbf{x})$ of the entry condition (4), we restrict our attention to cases when $n^\star(\mathbf{x})$ is continuous with finite first order derivatives. We denote by $n^\star(\mathbf{x})$ the maximal solution[18] to (4) for a given vector $\mathbf{x}$ of security investments. A preliminary result is that the solution $n^\star(\mathbf{x})$ of the Cournot equilibrium for the number of attackers is decreasing in x_i and namely

$$\frac{\partial n^\star(\mathbf{x})}{\partial x_i} < 0 \quad \text{where } n^\star(\mathbf{x}) \triangleq \max\{0, \arg\operatorname*{solve}_n\{n - \tfrac{1}{N}\sum\nolimits_{i=1}^{N} \rho_i \cdot \sigma_i(x_i, n) = 0\}\} \tag{5}$$

[15] In the equilibria considered in this paper, insured targets will only choose to pay the security expenditure mandated by their insurer. Hence, we abuse notation slightly be using x_i to denote both the mandated expenditure and the actual expenditure.

[16] In the case of a successful attack that results a visible loss, digital forensics experts are recruited to investigate (Marcella Jr and Greenfield, 2002). Therefore, if a firm committed to set of security investments in a contract, an insurer can verify their realization either before a security event (Olakunnle, 2014) or after it (Werlinger et al., 2010).

[17] This assumption is standard in the insurance literature but less obvious for cybersecurity. The recent initiatives on data sharing by the EU Commission (see the cited NIS Directive) and the US federal government (`https://www.dhs.gov/sites/default/files/publications/Overcoming%20Perceived%20Obstacles%20White%20Paper_1.pdf`), as well as community initiatives such as the VERIS Database (`http://veriscommunity.net`) are likely to resolve this issue in the upcoming years.

[18] A natural assumption is a preference for attackers to enter the market. Therefore the realized configuration in practice would be the one with the largest number of attackers.

where $\arg\operatorname{solve}\{.\}$ returns the argument that is the solution to a non-unary function in terms of the function's other parameters. This follows directly from Assumption $\mathcal{A}.5$ by taking the partial derivative of both side of (5) and observing that $\mathcal{A}.5$ implies $N \geq \sum_{j=1}^{N} \rho_j \frac{\partial}{\partial n}\sigma_j(x_j, n)\big|_{n=n^\star(\mathbf{x})}$.

By assumption, attackers are in a Cournot subgame, and the Nash equilibrium is defined over the set of N target firm strategies as the simultaneous solution of all targets i of the following problem

$$\max_{x_i} \mathbb{E}[U_i(x_i)|n] \qquad \text{subject to} \qquad x_i \geq 0 \text{ and } n = n^\star(\mathbf{x}).$$

in addition the target optimization is subject to problem specific constraints generated by the specific functional form of $U_i(\cdot)$ and $\sigma_i(\cdots)$ chosen.

This is solved by setting the usual first order condition[19] $\partial\mathbb{E}[U_i(x_i)|n]/\partial x_i = 0$, substituting $n = n^\star(\mathbf{x})$ in the resulting derivative, and solving for x_i. The simultaneous solution for every target yields a number of points $\mathbf{x}^*$ that solve the first order condition under the constraint $n^* = n^\star(\mathbf{x}^*)$.

The fist order condition for the expected utility of targets is expanded as follows:

$$\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i} = -\mathbb{E}[U_i'(x_i)|n] - L_i\frac{\partial\sigma_i(x_i,n)}{\partial x_i}U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i)). \tag{6}$$

This decomposition contains a positive term and a negative term which capture the interplay between the risk aversion of the target and the marginal effectiveness of self-protection expenditures. The first term is negative and captures the unwillingness of the target to increase its spending to counter marginal increases in risk. The second term increases along the effectiveness of security expenditures, and is amplified by marginal utility of the risk averse target at the point of the maximally insurable loss.

To illustrate the properties of the model, we first consider the case of the risk neutral target. If target i is risk neutral its expected utility is simply the expected net monetary value of its assets $\mathbb{E}[x_i|n] = W_i - x_i - \sigma_i(x_i, n)L_i$. In the insurance literature the probability of a negative event is usually only affected by the type of the individual (e.g. Einav et al. (2013)) or exogenous parameters (e.g. Dionne and Wang (2013)) In our scenario, the probability σ_i is partly determined by the strategic effort of the agent (x_i) and this by itself has a major impact. Since the function σ_i is convex in the security investment argument ($\mathcal{A}.3$ and $\mathcal{A}.4$) the expected utility of the risk neutral target is no longer a straight line which would have made the optimal value of security investment to $\mathbf{0}$. The optimal level of defensive expenditure is obtained by simultaneously setting to zero for all target $i = 1$ the usual first order condition:

$$\partial\sigma_i(x_i, n)/\partial x_i = -1/L_i \quad \text{where } n = n^\star(\mathbf{x}) \tag{7}$$

Given the assumptions on σ a solution for the Nash equilibrium, $\mathbf{x}^\sharp \geq \mathbf{0}$ always exists and is unique. The expenditure $x_i^\sharp$ also represents the overall optimal expenditure for unregulated risk-averse target in presence of actuarially fair insurance. We discuss this issue in Section 3.1.

To show existence and unicity of a Nash equilibrium one could follow Theorem 8 of Rosen (1965) and show that the entry condition in (1.2) describes a convex set of admissible

[19] In the general case where externalities are explicitly considered both in the utility functions and in the probability of a successful attacks, the general solution would need to solve the equation $\operatorname{diag}(\nabla\mathbb{E}[\mathbf{u}(\mathbf{x})|\mathbf{x}, n]) = \mathbf{0}$ subject to $n = n^\star(\mathbf{x})$, where $\mathbf{0}$ is an N length null vector and $\mathbf{u}$ is the vector of all utility functions $\mathbf{u} = [U_i]_{i\in\{1,\ldots,N\}}$. In our model $\partial\mathbb{E}[U_i(x_i)|n]/\partial x_j] = 0$, for all $j \neq i$ by assumption as all externalities are captured by the n term. At the equilibrium n must be replaced by $n^\star(\mathbf{x})$ and thus $\partial/\partial x_j(\,\partial\mathbb{E}[U_i(x_i)|n]/\partial x_i|_{n=n^\star(\mathbf{x})})$ is in general different from zero.

points and restrict the analysis to models with global diagonal strict concavity as the latter guarantees that any sequence of optimizing changes starting in $\mathbf{x}$ would eventually lead to the attractor $\mathbf{x}^*$ corresponding to the Nash equilibrium[20]. However, the diagonal strict concavity assumption may be overly restrictive as several interesting cases exist for which it does not hold[21].

In practice, targets are unlikely to start from a very high level of security expenditures close to their upper bound L_i. Hence, we restrict our analysis to a bounded subset around the origin including the first solution (if any) to the Nash equilibrium i.e. for $x_i \in [0, \bar{x_i}]$ such that $x_i^* \leq \bar{x}_i$ where x_i^* is the smallest solution to (6). In this case a simultaneous Nash equilibrium always holds by setting three conditions on the elasticity of the drop in utility, $\mathrm{e}_{\Delta U_i}$.

Proposition 2.1
A solution to the first-order condition $\mathbf{x}^* \in [\mathbf{0}, \bar{\mathbf{x}}]$ is a unique and stable Nash equilibrium in the region $[\mathbf{0}, \bar{\mathbf{x}}]$ if (i) the elasticity of the drop in utility is smaller than the elasticity of the marginal increase in successful attack by changing the number of attackers, $\mathrm{e}_{\Delta U_i} \leq -\mathrm{e}_{\partial \sigma_i / \partial n}$, (ii) it is also smaller than half of the elasticity of the marginal increase in successful attack by changing security expenditures, $2\mathrm{e}_{\Delta U_i} \leq -\mathrm{e}_{\partial \sigma_i / \partial x_i}$, and (iii) in the absence of security investments, $\mathbf{x} = \mathbf{0}$, $n_0 = n^\star(\mathbf{0})$, it must be smaller than the combination of the elasticity of probability of a successful defense and the marginal utility discounted by the expected drop in utility $\mathrm{e}_{\Delta U_i(0)} \leq -\mathrm{e}_{\sigma_i(0,n_0)} - U'(W_i)/(\sigma_i(0, n_0)\Delta U_i(0))$.

Under our general assumptions $\mathcal{A}.1$–$\mathcal{A}.7$ and the above conditions (i) and (ii) in the region $[\mathbf{0}, \bar{\mathbf{x}}]$, the marginal utility of the targets evaluated at $n = n^\star(\mathbf{x})$ is a monotone decreasing function on the constrained surface determined by the attackers' entry condition. Condition (iii) implies that the function that has an initial positive point at $\mathbf{x} = \mathbf{0}$ and therefore $\mathbf{x}^* \in [\mathbf{0}, \bar{\mathbf{x}}]$ is a maximum.

All three conditions compare the elasticity of the gain derived from a security investment ($\mathrm{e}_{\Delta U_i}$) with the ability of the same target to obtain such a gain by investing x_i. The relative weight of condition (i) and (ii) tells us that in comparison to the potential gap from security expenditures ($\mathrm{e}_{\Delta U_i}$) the direct impact of security expenditures ($\mathrm{e}_{\partial \sigma_i / \partial x_i}$) can be half the size of the the potential indirect impact ($\mathrm{e}_{\partial \sigma_i / \partial n}$). Since $\mathrm{e}_{\Delta U_i(x_i)} > 0$ for a risk averse target (from $\mathcal{A}.1$), condition (i) is possible from assumption $\mathcal{A}.2$ and the second part of assumption $\mathcal{A}.3$ which guarantees that $\mathrm{e}_{\partial \sigma_i / \partial n} < 0$. The possibility of condition (ii) to be true rests on assumptions $\mathcal{A}.2$ and $\mathcal{A}.4$ which make $\mathrm{e}_{\partial \sigma_i / \partial x_i} < 0$.

Condition (iii) requires that the benefit of action must be greater than the ability of transforming security investments into smaller chances of attacks, captured by $-\mathrm{e}_{\sigma_i}$ and the marginal utility of inaction and the resistance of a target to spend, captured by $U_i'(0)$. Inaction includes the risk of a drop in utility by itself and hence the last term must be discounted by such expected drop. Condition (iii) is possible due to assumptions $\mathcal{A}.3$ and $\mathcal{A}.1$.

If the conditions above are not satisfied by the risk averse utility function of interest, there might not necessarily be a unique Nash equilibrium[22]

[20]Consider the dynamic system where $\mathbf{x}(0)$ is a random initial endowment in $[\mathbf{0}, \bar{\mathbf{x}}]$. Setting $n(t+1) = n^\star(\mathbf{x}(t))$, the update in firm investment is denoted $\mathbf{x}(t+1) = [\arg\max_{x_i} \mathbb{E}\left[U_i(x_i)|n\right]_{n=n(t+1)}]_{i,\ldots,N}$. A vector $\mathbf{x}^*$ is a unique attractor if for any starting point in $[\mathbf{0}, \bar{\mathbf{x}}]$ the vector $\mathbf{x}^*$ is the unique point of convergence of the iterative sequence.

[21]For a negative exponential CARA utility functions diagonal strict concavity always holds. It does not hold, in general, for iso-elastic or power utility functions where agents exhibit CRRA type preferences.

[22]For identical CARA targets with a constant absolute risk aversion γ_i and constant elasticity α_i of the

2.2. A Benevolent Social Planner Mandating Security Investments

Before introducing the insurance market, it is useful to ascertain the optimal investment policy that a fully informed benevolent social planner would mandate. We will now show that in the presence of a public policy acting as a benevolent social planner, the socially optimal level of investment $x_i^\dagger$ will be greater than x_i^*.

By "benevolent" we adhere to a classical utilitarian definition where the social planner's utility function respects the preferences of the population of targets. This can be captured by the aggregate von Neumann–Morgenstern utility function which sums the utilities of the individual targets weighted by the values ν_i assigned by the policy maker to different targets. The planner's action is to "mandate" security investments for each target denoted $x_i^\dagger$ and we will assume that $x_i^\dagger$ is binding, measurable, and enforceable (compulsory security standards).

Hence, the expected utility of the policy maker for all security expenditures $\mathbf{x}$ and number of attackers n is the linear combination of the expected utilities of the individual targets $\mathbb{E}[U_i(x_i)|n]$ as specified in (3).

$$\mathbb{E}[U_P(\mathbf{x})|n] = \sum\nolimits_{i=1}^{N} \nu_i \, \mathbb{E}[U_i(x_i)|n]. \tag{8}$$

We model the outcome as a subgame-perfect equilibrium. The choices of the policy maker in the first stage of the game must be optimal given the strategies of the players in the second stage. The strategy of each actor in the second stage must also be optimal for each possible security investment of the policy maker in the first stage and given the strategies of all other second-stage players, Binmore (2007).

In stage two, targets are not active players in the game without insurance because the policy maker mandates their level of defensive expenditure, x_i. Still, the payoffs of the targets are important since the policy maker chooses the levels of defensive expenditure to maximize their expected utility (8).

Each potential attacker can still choose whether or not to participate in attacks against the population of targets. For potential attackers to be part of a subgame-perfect equilibrium, it is sufficient for the equilibrium number of attackers per target, $n^\star(\mathbf{x})$, to satisfy (4) for each set of feasible defensive expenditures $\mathbf{x}$. Hence, they are indifferent between participating or not participating in attacks.

Since the equilibrium number of attackers per target, n^*, adjusts to changes in the levels of defensive expenditure, the policy maker must take this adjustment into account when determining the optimal choice in the first stage of each game. Therefore, the optimal choice of the policy maker satisfies the usual first-order conditions: $\partial\mathbb{E}[U_P|n^\star(\mathbf{x})]/\partial x_i = 0$ for all i under the constraint represented by the attacker indifference condition (4). The Nash equilibrium point $(\mathbf{x}^*, n^*)$ of the unregulated targets is just one of the evaluated expected utilities for the policy maker since at the equilibrium, hence we can see that $n^* = n^\star(\mathbf{x}^*)$.

probability of a successful attacks w.r.t. absolute and marginal security expenditures the first two conditions are (i) $\gamma_i < \alpha_i$, (ii) $\gamma_i < 1/2\alpha_i$ whilst condition (iii) converges asymptotically to (i) for large losses L_i. As soon as the risk aversion of the target is less than half the effectiveness of security expenditures, a unique Nash equilibrium will exist. For identical targets with a CRRA utility function, a small coefficient for relative risk aversion $\xi_i < 1$, and a constant or non-increasing ratio between the loss and the wealth as the wealth W_i becomes bigger, condition (iii) converges to $\alpha_i > 0$. Conditions (i) and (ii) will also be verified for the entire feasible range. Hence for large wealths W_i, a unique equilibrium exists, irrespective of the effectiveness of defenses α_i. If the target has a relatively high coefficient of relative risk aversion $\xi > 1$, a unique Nash equilibrium might not exist: at an one equilibrium, targets starting from high initial security expenditure will maintain such expenditure and push attackers out of the market; in another equilibrium, targets will have low security expenditures and a large number of attackers. A computer program that simulates this effect is available in the supplementary materials.

The decomposition of the marginal expected utility illustrates how incentives of policy makers may differ from incentives of individual unregulated targets[23]:

$$\frac{\partial \mathbb{E}[U_P(\mathbf{x})|n^\star(\mathbf{x})]}{\partial x_i} = \underbrace{\nu_i \frac{\partial \mathbb{E}\left[U_i(x_i)|n\right]}{\partial x_i}\bigg|_{n=n^\star(\mathbf{x})}}_{\text{Unregulated Nash Equilibrium}} + \overbrace{\frac{\partial n^\star(\mathbf{x})}{\partial x_i}}^{\text{Drop in attackers if } dx_i>0} \cdot \underbrace{\frac{\partial \mathbb{E}\left[U_P(\mathbf{x})|n\right]}{\partial n}\bigg|_{n=n^\star(\mathbf{x})}}_{\text{Global benefit if } dn<0} \tag{9}$$

The divergence between the policy maker's marginal utility and the individual target's marginal utility occurs because individual targets ignore the beneficial effect that the target's expenditure has in reducing attacks on other targets: the number of interested attackers decreases as target i strengthens its defenses ($-\partial n^\star(\mathbf{x})/\partial x_i \geq 0$ and decreasing in x_i) and the diminishing number of attackers increases the marginal expected utilities of the policy maker – i.e. the aggregated utilities of all targets ($\partial \mathbb{E}\left[U_P(\mathbf{x})|n\right]/\partial n \leq 0$). The overall term is positive and shift to the right of the optimal choice of the security expenditure.

Proposition 2.2
The security investment $x_i^\dagger$ mandated by a benevolent social planner to a risk averse target i is larger than the optimal security investment x_i^* that the same target i would have chosen in an unregulated environment, i.e. $x_i^\dagger \geq x_i^*$.

That the benevolent social planner mandating investments can "improve", strictly in an individual welfare sense, the outcomes for all targets is a well-understood effect. In the absence of the social planner and when the number of targets is large, no single target can, by altruistically and unilaterally raising their defensive expenditure, can reduce the overall number of attackers n_i^*. Hence, their unilateral increase has only the effect of shifting them from their optimal expenditure. From Assumptions $\mathcal{A}.1$ to $\mathcal{A}.7$ for a given n_i^*, the optimal expenditure x_i^* is unique and all deviations from this expenditure are sub-optimal. In contrast, the social planner accounts for the attacker reaction and by mandating expenditure across *all* targets attains a higher overall utility for each target than each individual target could do by acting alone.

3. Cyberinsurance Contracts

We now introduce an insurance market that provides targets with coverage against losses from cyber-attack. At first we assume a perfectly competitive market providing actuarially fair insurance. Then we move on to the opposite case when a single monopoly insurer can extract a full surplus from targets.

3.1. Actuarially fair cyber-insurer and unregulated targets

In the absence of the social planner individual targets can freely choose a level of defensive expenditure x_i as well as whether to purchase an insurance contract specified by the tuple (q_i, ℓ_i). As such the i target's expected utility is

$$\mathbb{E}[U_i(q_i, \ell_i, x_i)|n] = \sigma_i(x_i, n)U_i(W_i - x_i - q_i - \ell_i) + (1 - \sigma_i(x_i, n))U_i(W_i - x_i - q_i) \tag{10}$$

[23]The second term of the policy maker marginal utility $\partial \mathbb{E}\left[U_P(\mathbf{x})|n\right]/\partial n$ can also be expressed as the ν_j-weighted sum of the product of the impact of loss on utility $U_j'(W_j - x_j - \mathfrak{L}(W_j - x_j, L_j))$ for the decrease of successful attacks if $dn > 0$ represented by $-\partial\sigma_j(x_j, n)L_j/\partial n$. This decomposition is more amenable to quantification.

As the quantity $\sigma_i(x_i, n_i)$ is the probability that target i incurs a loss from an attack, it also represents the probability that an insurer who insures target i for contract (q_i, ℓ_i) will pay out the amount $L_i - \ell_i$. Since insurance markets are efficient, the insurer does not make a profit at the equilibrium:

$$q_i = \sigma_i(x_i, n)(L_i - \ell_i) \tag{11}$$

This condition is commonly referred to as actuarially fair insurance. Consider now n as fixed exogenously. Target i wishes to choose x_i and (q_i, ℓ_i) to maximize the expected utility $\mathbb{E}[U_i(x_i, q_i)|n]$. Risk neutral targets are indifferent between buying insurance coverage or "self-insuring", even when this coverage is provided at an actuarially fair price. Hence, they always choose no insurance ($\ell_i = L_i$ and $q_i = 0$).

When targets are risk averse, the analysis requires a more careful reasoning. At first we consider again a noncooperative game among the players. The choices of attackers are unchanged w.r.t. their choices in absence of efficient insurance markets (§2). In the absence of insurance, the strategy for a target at the equilibrium was simply the choice of defensive expenditure. When targets can also purchase actuarially fair insurance, target i's strategy involves two choices: (i) the level of deductibles ℓ_i and (ii) the level of defensive expenditure, x_i. The premium is then determined by (11) given our assumption about the efficiency of insurance markets.

Lemma 3.1
For a given number n of attackers per target, a risk averse target i which is offered insurance at an actuarially fair rate will always find it optimal to choose a level of coverage equal to the full loss ($\ell_i = 0$) and select the same level of security expenditures of a risk neutral targets ($x_i = x_i^\sharp$).

The general solution of the Nash equilibrium requires the simultaneous solution of setting to zero of the partial derivative of the expected value in (10) with respect to x_i and ℓ_i under the free entry constraint represented by (4). When target i is risk averse and $U_i(w)$ is weakly concave, it is convenient to solve for target i's optimal choice in two steps. At first, we calculate the optimal deductible $\ell_i(x_i)$, for each level of defensive expenditure x_i. Then, we calculate the optimal level of defensive expenditure x_i when ℓ_i is set to its optimal level, that is, when $\ell_i(x_i)$ is substituted for ℓ_i in (10).

The choice to be fully insured ($\ell_i = 0$) follows directly by applying Jensen's inequality to the expanded version of (10). The expected utility will then reduce to a single state as the quote is payable in both periods and the coverage is complete as there is no deductible. This means that the i target will choose x_i to maximize the utility $U_i(W_i - x_i - \sigma_i(x_i, n_i)L_i)$. Since $U_i(w)$ is an increasing function, this corresponds to choosing x_i to maximize the expected net value of target i's assets, and namely $W_i - x_i - \sigma_i(x_i, n_i)L_i$. Unsurprisingly, a risk averse target who is able to offload the entire risk of a loss by the purchase of actuarially fair insurance chooses the same level of defensive expenditure as would be chosen by a risk-neutral target as determined by (7).

Note that the target that purchases actuarially fair insurance will spend more than the target that opts for self-protection: for a given number of attackers per target n, the former will spend $x_i^\sharp - \sigma_i(x_i^\sharp, n)$, whereas the latter will spend $x_i^\sharp$. The relation between $x_i^\sharp$ and the security expenditure of the risk neutral target x_i^* depends on the actual shapes of U_i and σ_i but it is possible to establish some general result for risk averse targets. The following decomposition clarifies the relations between the expected utility of the risk averse target

and the expected utility of the target when actuarially fair insurance is available.

$$\overbrace{\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}}^{\text{Ex. Util. at Nash}} = \overbrace{\frac{\partial U_i(\mathbb{E}[x_i|n])}{\partial x_i}}^{\text{Ex. Util. Fair Ins.}} + \overbrace{U'(\mathbb{E}[x_i|n]) - \mathbb{E}[U'(x_i)|n]}^{\text{Degree of Prudence}} + \underbrace{-\frac{\partial \sigma(x_i,n)}{\partial x_i} L_i}_{\text{Marginal Loss in } x_i} \cdot \underbrace{\left(U'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i)) - U'(\mathbb{E}[x_i|n])\right)}_{\text{Marginal Util. at Local Risk Neutrality vs Expected } x_i} \tag{12}$$

If the first two terms are positive, then at the optimal value of investments for the risk neutral target $x = x^\sharp$ the term relating to the actuarially fair insurance converges to zero, but the derivative of the expected utility at the Nash equilibrium it is still positive and the maxima is farther away; hence $x^* \geq x^\sharp$.

The second term of the decomposition captures the degree of prudence of the target, that is the sign of U_i''' which for prudence is set to be $U_i''' < 0$ and imprudence in the opposite case. In this case the reverse of Jensen's inequality operates as U_i' is concave, hence $U'(\mathbb{E}[x_i|n]) \leq \mathbb{E}[U'(x_i)|n]$ so this term is negative. However, the larger the concavity of U_i', the steeper the monotone decrease in U_i' will be given $U_i'(w_1) \geq U_i'(w_2)$ for $w_1 \leq w_2$. As such, the maximally insurable loss $W_i - x_i - \mathfrak{L}(W_i - x_i, L_i)$ will be likely reached in close proximity with the point $W_i - L_i$ and will be far closer to $W_i - L_i$ than the point $\mathbb{E}[x_i|n] = W_i - x_i - \sigma(x_i, n)L_i$. Therefore the third term can (and will in most plausible cases) be positive and substantially larger than the second term, this leads us to the following result.

Theorem 3.2
Let $x_i^\sharp$ be the optimal security investment for the risk averse target with actuarially fair insurance premium $q_i^\sharp = \sigma_i(x_i^\sharp, n_i)L_i$ and x_i^* the optimal security investment in the absence of any available insurance contracts. The inequality $x_i^* > x_i^\sharp$ occurs if (i) the elasticity of the drop in utility is smaller than the elasticity of the marginal success of security expenditures in the number of attackers $\mathrm{e}_{\Delta U_i(x_i)} < -\mathrm{e}_{\partial \sigma_i(x_i,n)/\partial n}$ for all values of $x \leq x_i^\sharp$ and $n = n^\star(\mathbf{x})$ and (ii) the expected marginal utility at $x_i^\sharp$ is smaller than the marginal utility at the maximally insurable loss at $W_i - x_i^\sharp$, $\mathbb{E}[U_i'(x_i^\sharp)|n^\sharp] \leq U_i'(W_i - x_i^\sharp - \mathfrak{L}(W_i - x_i^\sharp, L_i))$.

The first condition implies that for $x_i < x_i^\sharp$ the marginal effectiveness of security expenditures to limit the number of attackers is greater than the aversion of the target to additional deterministic expenditure. The interplay between σ_i and U_i can change the relative position of the optimal security expenditure for uninsured versus insured targets and is represented by condition (ii).[24] We can see that for uninsured targets the optimal level of expenditure x_i^* bounds the marginal expected loss as follows: $-1/\lambda_i \leq L_i \, \partial \sigma_i(x_i, n)/\partial x_i|_{x_i = x_i^*} \leq -\lambda_i$. At the same time by (7) we recover $-1/\lambda_i < -1 = L_i \partial \sigma_i(x_i, n)/\partial x_i|_{x_i = x_i^\sharp} \leq -\lambda_i$. Therefore the marginal chances of a successful attacks for the optimal investment $x^\sharp$ of the insured target is essentially in the same narrow interval as the marginal chances for the optimal investment x_i^* for the uninsured target. The precise crossover point requires us to specify a functional form for σ_i and U_i; however, this crossover point will exist given the underlying assumptions.

[24] For targets with a exponential CARA utility with a constant for absolute risk aversion γ and a constant elasticity α of success probability w.r.t. absolute and marginal expenditures condition (i) is equivalent to $\gamma < \alpha$ and condition (ii) converges to (i) for large L. When actuarially fair insurance is available, the optimal investment coincides with Nash equilibrium case for risk neutral targets.

The simultaneity assumption on the choices of attackers and targets are implies that an individual target neglects the effect that a change in the target's level of defensive expenditure has on the incentives of potential attackers to mount attacks. Furthermore, each potential attacker is assumed to neglect the effect that the attacker's decision might have on the overall level of threat perceived by the targets and on the targets' levels of defensive expenditure. These assumptions appear to be plausible approximations when the number of potential attackers and the number of targets is large. For in this case, a change in the choice of defensive expenditure by a single target is not likely to affect the overall expected reward from attacks by very much. Similarly, a change in the participation decision of a single potential attacker is not likely to have a significant effect on the number of attackers per target.

3.2. Policy Maker with Efficient Insurance Markets

To our efficient insurance markets case, we now introduce a benevolent social planner with the same utilitarian objectives as described in Section 2.2. The policy maker, the targets, the attackers and the insurers satisfy the same assumptions and objective functions of the previous sections. We again model the situation as a two stage game with a subgame-perfect equilibrium.

Under hypothesis of efficient insurance markets, insurer's profits are zero and we can ignore this component in the utility function of the policy maker. Therefore, the objective function of the policy maker is determined by (8) where the utility of the targets is determined by (10) in place of (3) from Section 2.

The policy maker still chooses the level of defensive expenditure for each target in stage 1. In stage 2, each potential attacker also chooses whether or not to participate in attacks. The choices available to the targets are instead different from the no insurance case because each target i can decides the insurance contract, (q_i, ℓ_i). Targets choose a level of coverage to maximize the expected utility given in (10) where $x_i^\dagger$ is determined exogenously by the policy maker. If the insurance industry provide actuarially fair insurance q_i will be determined by (11).

A strategy for the policy maker is simply the policy maker's choice of defensive expenditures for the targets. Strategies for stage 2 players are more complex. A strategy for each potential attacker is conditioned on whether or not to participate in attacks for each possible set of choices by the policy maker in the first stage. All strategies for a target are conditioned on the chosen level of coverage for each possible choice by the policy maker in the first stage.

In a subgame-perfect equilibrium the chosen ℓ_i, for each target i must be optimal for each set of defensive expenditures that could be chosen by the policy maker when the number of attackers per target is also given by the equilibrium level n_i^*. If a target is risk neutral, $\ell_i = L_i$ is, of course, always optimal. For the case where a target is strictly risk averse, Proposition 3.1 asserts that $\ell_i = 0$ is optimal for all levels of defensive expenditure x_i and all levels of n_i. In both cases, each target i always receives the expected utility value of its assets, $U(W_i - x_i - \sigma_i L_i)$.

Since the policy maker anticipates this outcome in stage one of the game, the policy maker's expected utility from (8) can be rewritten as

$$\mathbb{E}[U_P(\mathbf{x})|n] = \sum\nolimits_{i=1}^{N} \nu_i U_i(\mathbb{E}[x_i|n]) \text{ with } n = n^\star(\mathbf{x}). \tag{13}$$

The function $n^\star(\mathbf{x})$ has also been substituted for n in the expression for $\mathbb{E}[U_i]$ because the policy maker forecasts the response of potential attackers to different levels of defensive expenditure in the second stage of the game.

The equilibrium levels of defensive expenditure determined in the first stage of the game are assumed to satisfy the usual first order conditions for optimality: $\partial\mathbb{E}[U_P]/\partial x_i = 0$ for all i. Similarly to (9) for the regulated targets in the absence of insurance, this condition can be decomposed into two components that clarify why the investment mandated by the policy maker is larger than the investment by unregulated targets, even in presence of an efficient cyberinsurance market.

$$\frac{\partial\mathbb{E}[U_P(\mathbf{x})|n]}{\partial x_i} = \nu_i \underbrace{\left.\frac{\partial U_i(\mathbb{E}[x_i|n])}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}}_{\text{Risk neutral optim. when=0}} + \overbrace{\frac{\partial n^\star(\mathbf{x})}{\partial x_i}}^{\text{Drop in attackers if } dx_i>0} \cdot \underbrace{\left.\frac{\partial U_P(\mathbb{E}[\mathbf{x}|n])}{\partial n}\right|_{n=n^\star(\mathbf{x})}}_{\text{Global benefit if } dn<0}. \tag{14}$$

Once again, each target ignores the beneficial effect the target's expenditure has in reducing the number of attackers not only on itself but on the other targets as well. This phenomenon is captured by the second term of the decomposition which is the pro quota variation to the social expectation of a reduction in the number of attackers ($\partial\mathbb{E}[U_P]/\partial n$) due to the reduction of this very number of attackers thanks to the increase in expenditure by the i-th target $\partial n^\star(\mathbf{x})/\partial x_i$.

Proposition 3.3
In presence of efficient markets for cyberinsurance the security investment $x_i^\ddagger$ of a risk averse target i mandated by a benevolent policy maker to each individual target is larger than the optimal security investment $x_i^\sharp$ that the same target i would choose in an unregulated environment with actuarially fair insurance ($x_i^\ddagger \geq x^\sharp$).

3.3. Mandated Protection from Monopolist Cyber Insurer

Under Assumptions $\mathcal{A}.9$ and $\mathcal{A}.10$ an insurer can mandate a minimal level of defensive expenditure as part of the insurance contract conditions and compute quotes for each target under this investment assumption. This scenario ought to be the best chance for the incentives of a profit-maximizing insurer to align with those of a "benevolent" policy maker. In contrast, we show that the choice of defensive expenditure by the insurer may not be socially optimal and even worse that self-protection as individual security expenditures collapse and quotes increase.

The outcome of the interaction between the insurer and the targets is modeled as a subgame-perfect equilibrium of a two-stage game. In the first stage of the game, the insurer makes its offer to each target. The offer consists now of a triple (q_i, ℓ_i, x_i) where the first two arguments are respectively the premium and the deductible as in the previous section and x_i is the required level of defensive expenditure which an insured target must incur. The total profit obtained by an insurer would be the difference between the premium paid by the targets and expected losses that must be covered (minus the deductibles):

$$\Pi = \sum\nolimits_{i=1}^{N} q_i + \sigma_i(x_i, n^\star(\mathbf{x}))(\ell_i - L_i) \text{ where } \ell_i \leq L_i. \tag{15}$$

In the second stage of the game, targets and potential attackers make simultaneous choices. Each target must choose whether or not to accept the insurer's offer. If a target accepts, then no further choice is required. If a target rejects the offer, then the target must also choose the level of defensive expenditure x_i which it will incur. As in previous sections, each potential attacker must choose whether or not to participate in attacks on the targets.

The monopolist wishes to offer an insurance contract which all targets will be willing to purchase given the overall security environment consequent to all targets choosing insurance;

recalling that the level of risk is solely determined by the attacker entry condition $n = n^\star(\mathbf{x})$. For the targets to be willing to accept the insurer's offer, they must be indifferent between accepting or not accepting. Therefore, the following incentive compatibility constraint must hold:

$$\mathbb{E}[U_i(q_i, \ell_i, x_i)|n^\star(\mathbf{x})] \geq \mathbb{E}[U_i(x_i)|n^\star(\mathbf{x})]. \tag{16}$$

Given monopoly power, the single insurer

Since the insurer holds a monopoly will extract all possible surplus from each target and therefore we only consider the boundary form of the constraint.

The left-hand side of (16) is defined in (10) and denotes the expected utility which target i obtains by accepting the insurance contract. In an equilibrium, where all targets purchase insurance and choose the level of expenditure specified by the insurance contract $\mathbf{x_i}^\circ$, the appropriate forecast of the number of attackers is $n^\star(\mathbf{x_i}^\circ)$, the equilibrium number of attackers per target. The right-hand side of (16) indicates the expected utility which a target obtains by rejecting insurance ($q_i = 0$ and $\ell_i = L_i$) and making an optimal choice of defensive expenditure.

It is convenient to solve the insurer's problem in two steps. In Step 1, the insurer chooses $q_i(\mathbf{x})$ and $\ell_i(\mathbf{x})$ to produce the profit that satisfies the incentive compatibility constraint for each level of defensive expenditure, x_i. In Step 2, the insurer chooses the level of defensive expenditure that maximizes $\Pi(\mathbf{x})$.

When all x_i are held fixed, the right-hand side of (16) is a constant as is the quantity $\sigma_i(x_i, n^\star(\mathbf{x}))$ in the profit of the insurer from (15). The left-hand side of (16) is then smaller than target's utility when he chooses the level of deductible ℓ_i:

$$U_i(W_i - x_i - q_i(\mathbf{x})) \geq \mathbb{E}[U_i(q_i, \ell_i, x_i, n^\star(\mathbf{x}))] \tag{17}$$

The insurer is going to always offer the insuree a full insurance contract $(q_i, 0)$ as this would maximize the value of the offer for the target. The same reasoning applies when the target receives an offer for actuarially fair insurance as in Proposition.3.1. Since the insurer has a monopolist advantage the compatibility bound will hold as an equality. This provides an implicit expression for q_i and x_i:

$$U_i(W_i - x_i - q_i(\mathbf{x})) = \mathbb{E}[U_i(x_i)|n^\star(\mathbf{x}))]. \tag{18}$$

Since U_i is monotone increasing, and thus invertible, the implicit function above is well defined for $q_i(\mathbf{x})$ for all values of x_i and all parameters W_i, L_i.

Each q_i is a function of $\mathbf{x}$, the environment as a whole, and by extension the insurer must identify a vector function $\mathbf{q}(\mathbf{x})$ on the vector space $\mathbf{x}$. The Jacobian matrix $\nabla\mathbf{q}(\mathbf{x})$ captures how a change in mandated security expenditures $\mathbf{x}$ of targets affects simultaneously all insurance quotes (through the incentive compatibility constraint and the attackers' behavior from the Cournot equilibrium).

We introduce a term corresponding to the marginal ambient risk for the i-th risk averse target showing how the expected marginal loss changes as the number of attackers changes as discounted by the ratio in marginal utility when the target is asked to pay a given premium q_i w.r.t the maximally insurable loss $\mathfrak{L}(W_i - x_i)$.

$$\mathfrak{R}_i(x_i, q, n) = \overbrace{\frac{\partial\sigma_i(x_i, n)L_i}{\partial n}}^{\text{Raise Loss by } dn>0} \cdot \left(1 - \overbrace{\frac{U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i))}{U_i'(W_i - x_i - q))}}^{\text{Marg. Util. Max. Insurable Loss vs Quote}}\right) \tag{19}$$

The first term is always positive ($\mathcal{A}.2$). The sign of the second term depends from the level of risk aversion. Since U_i' is monotone decreasing ($\mathcal{A}.1$), if the quote is smaller than the maximally insurable loss ($q_i(\mathbf{x}) \leq \mathfrak{L}(W_i - x_i)$) then $U_i'(W_i - x_i - q_i(\mathbf{x})) \leq U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i))$ and therefore the marginal risk is negative. Increasing the expenditure x_i decreases the marginal environmental risk for quotes below the maximally insurable loss. For large quotes, the effect is reversed.

The marginal adjustment in the quotes q_j and q_i is given by the following equations for $i, j = 1 \dots n$ where we show how the quote q_i of the target i varies with its own mandated expenditures (the diagonal of the Jacobian) and the interacting terms q_j of the the other targets $j \neq i$ variate with x_i:

$$\frac{\partial q_i(\mathbf{x})}{\partial x_i} = -1 - \overbrace{\left.\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}}^{\text{Marg. Exp. Utility at Nash Eq. when =0}} \cdot \overbrace{\frac{1}{U_i'(W_i - x_i - q_i(\mathbf{x}))}}^{\text{Marg. Utility at Quote}} + \overbrace{\frac{\partial n^\star(\mathbf{x})}{\partial x_i}}^{\text{Drop attackers by } dx_i > 0} \cdot \overbrace{\mathfrak{R}_i(x_i, q_i(\mathbf{x}), n^\star(\mathbf{x}))}^{\text{Marg. Risk at Cournot Equil.}} \tag{20}$$

The interaction term contains the same components for the j target:

$$\frac{\partial q_j(\mathbf{x})}{\partial x_i} = \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \cdot \mathfrak{R}_j(x_j, q_j(\mathbf{x}), n^\star(\mathbf{x})). \tag{21}$$

The first two components of the functional form of $\partial q_i / \partial x_i$ are relatively obvious. The first term is a constant negative marginal rate showing that, ceteris paribus, the quote must decrease linearly with the mandated security expenditures x_i since the latter directly competes with the quote q_i for a direct share of the overall wealth $W_i - x_i - q_i$ of the target. The second term caters for the alternative solution of self-protection. For $x_i < x_i^*$ the marginal expected utility is positive and the overall term is negative as U_i' is always positive. Hence, for values of expenditures larger than the expenses at the Nash equilibrium, this term makes rising the quote less attractive in terms of upfront costs than raising the self-protection effort. When the security expenses gets larger than the value of the Nash equilibrium the marginal expected utility changes sign and the overall term gets positives so it creates the margin for an increase in the quote as the share of wealth received from the target. Such growth is reduced by a factor, the marginal utility of the target when he consider both the expenditure and the quote $U_i'(W_i - x_i - q_i(\mathbf{x}))$. Therefore an optimal value for q_i could be potentially be found for $x_i > x_i^*$, if the third term of the decomposition were identically zero.

The third term of (20) and the condition on the other partial derivative (21) captures the interaction with the environment. The impact of this term on the monopolists ability to indirectly control the risk environment is significant. Consider the case when $n^\star(\mathbf{x})$ is exogenous and fixed for all $\mathbf{x}$. The reservation utility of the target is only determined by the second term of (20) and hence fixed. The interacting term between i and j in (21) is identically zero. In this scenario, a monopolist insurer cannot affect the reservation utility of any target. The monopolist can only provide full insurance (at a quote that extracts the full surplus), mandating targets to invest in security at a risk neutral level.

The outcome in the general case when $n^\star(\mathbf{x})$ is endogenous is radically different. What is unexpected is the effect on the population of N targets choosing between insurance and self-protection when i diminishes his expenditure x_i. The unintended, but dramatically present,

effect is that target j is largely powerless to react when x_i diminishes its expenditure and makes the environment riskier. If the quote offered to the j^{th} target is above j^{th} own maximally insurable loss then the marginal risk is positive. Once we multiply it by the marginal change in the number of attackers, it becomes negative.

Hence j's option is to buy more insurance at higher premiums as this is the only component of the partial derivative and $\partial q_j/\partial x_i \cdot dx_i > 0$ for $dx_i < 0$. So, from the perspective of the insurance company, making i invest less in security has the overall beneficial effect of incentivizing all other parties j to accept a higher premium. For target i itself, a viable option is to buy more insurance and to spend less on security and hence contribute to an increase in the margin for the quote. As a result, the insurer has incentives to *lower* mandatory security expenditures to make insurance more appealing.

If the monopolist sets the initial insurance premium too high then targets will not diffuse to adopting full insurance. So the insurance company would need to initially specify some contract that ensures that targets are shifted to an equilibrium path (for example charging at actuarially fair rate $(\sigma_i(x_i^\sharp, n^\sharp)L_i, 0, x_i^\sharp)$, and then progressively raise the expenditure, whereby $(q_i^\circ, 0, x^\circ)$ is the outcome and that no individual target can do better by rejecting this offer. Only after all targets are in the neighborhood of $x^\sharp$ then the equilibrium will be maintained because any target is worse off by rejecting insurance. We illustrate this effect numerically in Section 4.

The insurer optimizes its profit by the usual first order condition on (15) via the implicit expression for $\mathbf{q}(\mathbf{x})$ from (18) and the attackers entry condition.

$$\frac{\partial \Pi(\mathbf{x})}{\partial x_i} = 0, \quad U_i(W_i - x_i - q_i(\mathbf{x}) = \mathbb{E}[U_i(x_i)|n^\star(\mathbf{x})]. \tag{22}$$

The first order condition can then be decomposed as follows:

$$\frac{\partial \Pi}{\partial x_i} = \overbrace{-1 - L_i \left.\frac{\partial \sigma(x_i, n)}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}}^{\text{Fair Insurance: Risk Neutral Exp. Utility}} - \overbrace{\left.\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}}^{\text{Self- protection: Ex. Util. at Nash}} \cdot \overbrace{\frac{1}{U_i'(W_i - x_i - q_i(\mathbf{x}))}}^{\text{Marg. Utility at Quote}} + \tag{23}$$

$$- \overbrace{\frac{\partial n^\star(\mathbf{x})}{\partial x_i}}^{\text{Drop attackers by } dx_i > 0} \cdot \overbrace{\sum_{j=1}^{N} \left.\frac{\partial \sigma_j(x_j, n) L_j}{\partial n}\right|_{n=n^\star(\mathbf{x})} \frac{U_j'(W_j - x_j - \mathfrak{L}(W_j - x_j))}{U_j'(W_j - x_j - q_j(\mathbf{x}))}}^{\text{Aggregate Change in Utility}}$$

The first term of the decomposition is the equilibrium condition for the risk neutral target (7). We know that it is monotone because of Assumptions $\mathcal{A}.3$ and $\mathcal{A}.4$, and it is larger than zero for $x_i < x_i^\sharp$. The second term is the marginal expected utility at the Nash equilibrium. We know it is greater that zero for $x_i < x_i^*$ if the elasticity conditions in Proposition 2.1 hold. The term on the last line captures the relative importance of the overall marginal ambient risk across all targets j. The aggregate term is always positive given Assumption $\mathcal{A}.1$ on the target marginal utility function and Assumption $\mathcal{A}.5$ and the consequent constraint on the marginal change in the number of attackers w.r.t. the increase in security expenditures, see (5). Theorem 3.2 provides the condition under which $x^\sharp \leq x^*$.

Theorem 3.4
The optimal security expenditure x_i^* chosen by the unregulated targets in the absence of any available insurance is larger than the expenditure x_i° mandated by a monopolist insurer if for every target i and all values of $x \leq x_i^\circ$, the elasticity conditions (i)-(iii) for the existence

of a Nash equilibrium from Proposition 2.1 are satisfied and (iv) the marginal probability of a loss at $x_i = x_i^\circ$ and $n = n^\star(\mathbf{x}^\circ)$ is smaller than $-1/L_i$ (namely $\partial\sigma(x_i, n)L_i\partial x_i \leq -1$).

The first three conditions captures the interplay between the effectiveness of the security expenditures and the risk aversion of the targets as discussed in Proposition 2.1. The last condition compares the very same effectiveness against a risk neutral investment. This requires a Monopolist insurer to choose a mandatory expenditure where the marginal effectiveness of each dollar spent on security measures is declined below the $-45°$ line that intersects the utility axes at the expected loss. The larger the losses, the lower the declination has to be. An easy corollary is that if the above conditions holds then it is also $x_i^\circ \leq x_i^\sharp$, as (iv) also implies that then first order condition for the risk neutral target has not reached zero.

If we eliminate the environmental interaction from the decomposition of the utility of the monopolist insurer, the second group of terms in (23), then the resulting decomposition would have explained the widely held belief, in particular among security experts, that mandating cyberinsurance would solve the cyber security problem by mandating an increase in the security spending. Unfortunately, that simpler decomposition would not account for combined effect of the global strategic interaction of the attackers, which can be attracted by an environment with lower security spending, and of the risk averse target behavior. In risky environments, the response of risk averse targets might not be the socially desirable one (spend more on security to make the environment less risky) but rather the most attractive one (buy off more insurance and let somebody else deal with the risk).

4. The Insurance Trap

We now graphically illustrate a prisoner's dilemma outcome to our preceding general results, which we refer to as the 'insurance trap' by imposing specific functional forms on $U_i(\cdot)$ and $\sigma_i(\cdot)$. Our first example uses homogenous firms with ex-ante identical losses and utility functions. In our second example we impose two types of firms, with different degrees of risk aversion and size of losses, to illustrate that are effects can be a) substantial in magnitude and b) have substantial differences in magnitude between small and large firms.

For both examples we consider targets with an exponential CARA utility $U_i(w_i) = -1/\gamma_i \exp(-\gamma_i w_i)$ and a probability of a successful attack with multiplicative components $\sigma_i(x_i, n) = \exp(-\alpha_i x_i) n_i^{\beta_i}$ where $\alpha_i > 0$ and $0 < \beta_i < 1$ are positive scalar parameters.[25] For simplicity of exposition we only consider cases where the optimal investments in security and the number of attackers per target respect the boundary constraint that $\ln n_i = \alpha_i/\beta_i x$, hence forcing $\sigma_i \in [0, 1]$.

4.1. Example 1: Illustration with homogeneous firms

When all targets are ex-ante identical the utility function is $U_i(w_i) = U(w)$, with absolute risk aversion $\gamma_i = \gamma$ and losses from a successful attack of $L_i = L$. Similarly, all firms are equally proficient in protecting their cyber assets hence $\sigma_i = \sigma$, with ex-ante identical scalar coefficients $\beta_i = \beta$ and $\alpha_i = \alpha$. In this set-up we can solve for the exact functional forms of the nash equilibrium in the absence of insurance or a policy coordinator, $x_i^* = x^*$ and for the cases of actuarially fair insurance $x_i^\sharp = x^\sharp$ and monopoly insurance $x_i^\circ = x^\circ$, when all firms choose, ex-post, identical actions. A full expansion of the model is presented in Appendix B.1.

[25] The Mathematica files used to compute these bounds are available in the supplementary materials.

Figure 1 presents what happens when actuarially fair insurance is introduced to a set of unregulated targets and Figure 2 illustrates the system's evolution when a monopolist insurance switches from the initial actuarially fair insurance offering to a profit-optimizing quote. In both figures, the ordinate axis shows the targets' utility function and the abscissa axis reports the level of wealth w_i prior to the realization of success or non-success of an attack on a target. For an insured target i this value is $W_i - x_i - q_i$, for an uninsured target j this reduces to $W_j - x_j$.[26]

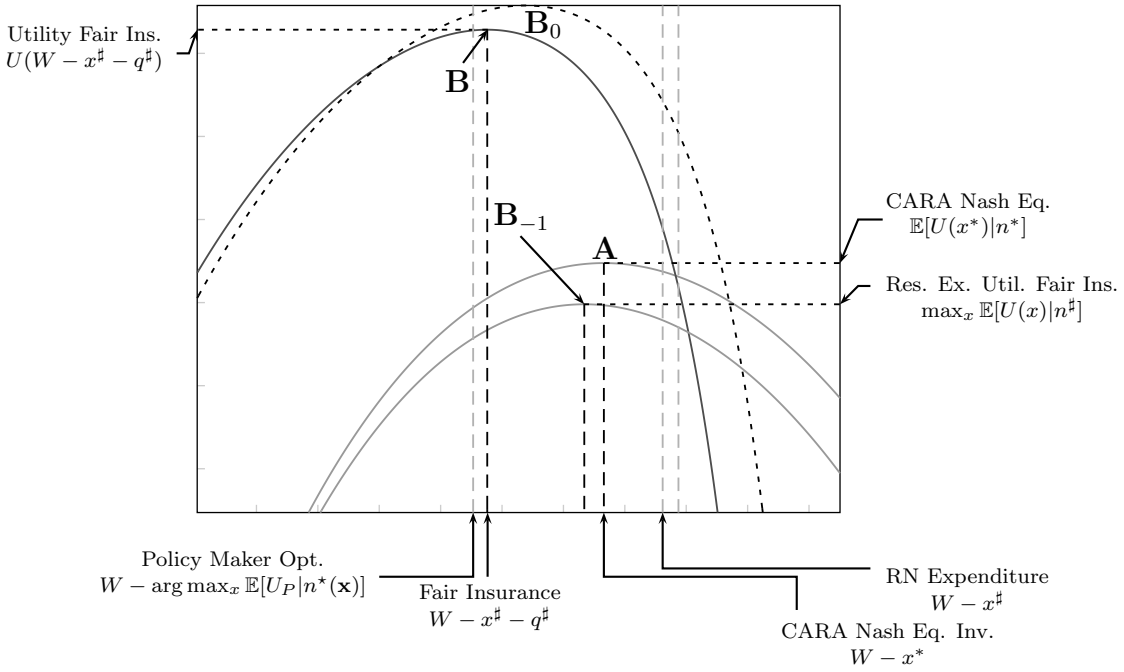

Figure 1: From Self-Protection to Actuarially Fair Insurance

The insurance trap for identical firms, the abscissa values run from zero to W, the 'wealth' of the identical firms. The ordinate axis is in utils, with lines emanating from the left axis representing certainty equivalence and those from the right expected utility. The expected utility at the simultaneous Nash equilibrium is labelled by point **A** when targets do not have access to any form of insurance. $\mathbf{B}_0$ is the certain utility of the first defectors when actuarially fair insurance is available (assuming the number of targets is large). **B** is the certain utility when all targets choose full coverage, at an actuarially fair price; with point $\mathbf{B}_{-1}$ the expected utility of a firm fully rejecting the actuarially fair contract.

In Figure 1, the sequence of points $\mathbf{A} \to \mathbf{B}_0 \to \mathbf{B} \not\to \mathbf{B}_{-1}$ that the targets will coalesce into clearly illustrates the mechanics of the insurance trap when actuarially fair insurance is introduced to unregulated, risk averse targets.

Point **A**, corresponding to the abscissa value x_i^*, is the optimal level of expenditure for each target under the simultaneous Nash equilibrium in the absence of insurance. Proposition 2.1 shows that a unique Nash equilibrium exists when the marginal rate of risk reduction α is greater than the level of constant absolute risk aversion, γ; the probability of a successful attack lies in the range $0 \leq \sigma \leq 1$ and each firm has a non-negative level of investment, $x_i \geq 0$.

[26]The curves here are plotted using the following parameters $\alpha_i = 0.035$, $\beta_i = 0.5$, $W_i = 100$, $L_i = 110$, $\gamma_i = 0.01$, $\rho_i = 2$. The viable domain for this model has a lower bound on the target investment $x_i \in [\beta_i/\alpha_i \log(\rho_i), L_i]$; outside this interval the number of attackers results in $\sigma_i > 1$.

After the Nash equilibrium **A** has been reached, we start to make available actuarially fair insurance to our targets. The presence of actuarially fair insurance would naturally result in risk averse targets wanting to fully insure, by setting their deductible to zero and adapt their expenditure relative to the equilibrium in the absence of insurance. When actuarially fair insurance is first provided, and all other targets investment is at $\mathbf{x}_{-i} = \mathbf{x}^*_{-i}$, an individual target can migrate cheaply to point $\mathbf{B}_0$ with a higher realized utility. However, all other targets will also migrate to full insurance, and the targets coalesce on the point **B** by paying the fair premium *and* the risk neutral (RN) expenditure $\mathbf{x} \to [x^\sharp]$. All targets are paying more than at Nash Equilibrium but each individual (hence the aggregate) security expenditure is lower than when fair insurance was not available ($x^*_i < x^\sharp_i$). Therefore, the number of attackers is higher in the presence of fair insurance ($n^* > n^\sharp$) as $n^\star(\mathbf{x})$ is monotonic decreasing in $\mathbf{x}$.

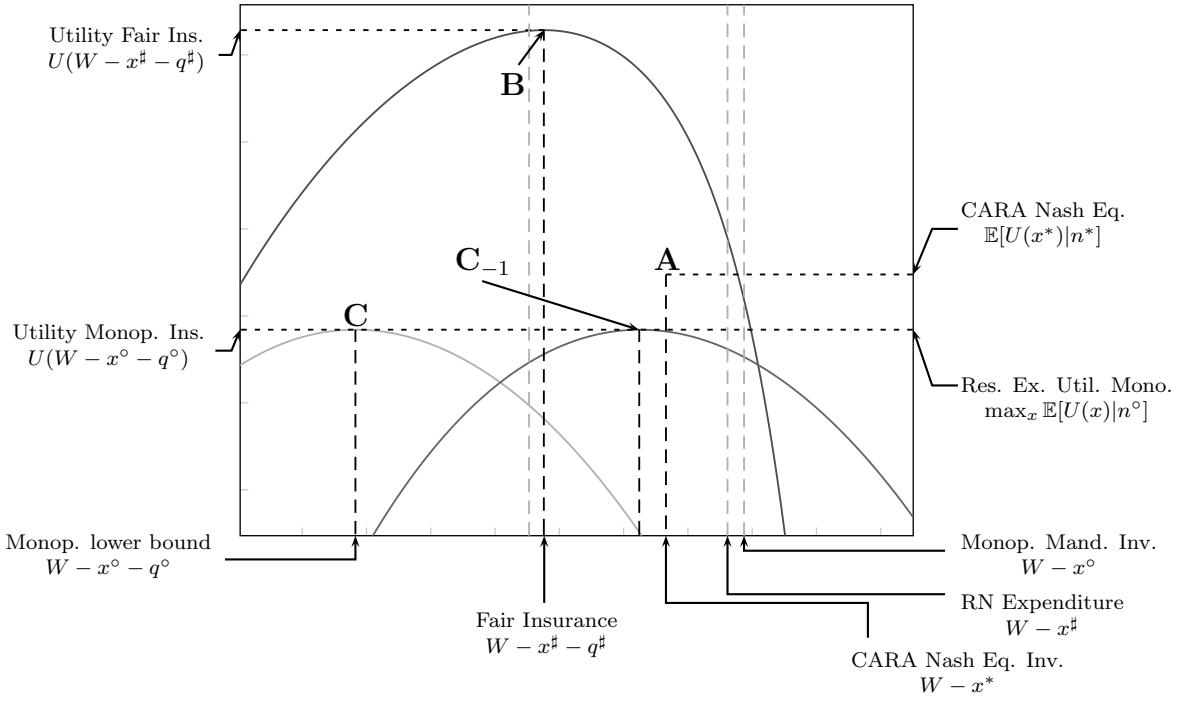

Figure 2: Geometry of the Insurance Trap

When targets have all migrated to full insurance, the expected utility at the Nash equilibrium, **A**, for uninsured targets is no longer available as attacking intensity is based on expenditure at **B**. The insurer can now adjust the expenditure and quote bundle until the target utility reaches point **C**, where the insurer attains the maximum payoff whilst ensuring that target certain utility is fractionally higher than point $\mathbf{C}_{-1}$, the reserve expected utility when the contract is rejected. Any insurance quote yielding certain utility above **A** for all targets allows the trap to be sprung.

Once all targets have migrated, the original optimal investment point, under the simultaneous Nash equilibrium with a maxima at x^*_i, no longer exists. A target wishing to unilaterally give up the available actuarially fair insurance would now shift to point $\mathbf{B}_{-1}$ as $\mathbf{x}_{-i} = [x^\sharp]$ and their investment will be $x^{\sharp\sharp} > x^* > x^\sharp$. Although their upfront expenditure would be lower than $x^\sharp + q^\sharp$, their expected utility will also be lower. Hence no firm would rationally reject the insurance contract. The insurer can then gradually shift the investment/insurance bundle to a higher risk higher premium combination.

Let us now assume that insurance is provided by a monopolist where the monopolist can mandate security expenditure as advocated by several academics and policy makers as we previously discussed. A monopolist insurer can now optimally choose to set target expenditure to $\mathbf{x} = [x_i^\circ]_i$ and set a vector of quotes $\mathbf{q} = [q_i^\circ]_i$, such that the target utility is at or above point $\mathbf{C}$. For any target unilaterally rejecting the insurance contract the alternative is to move to $\mathbf{C}_{-1}$ – the horizontal line represents the incentive compatibility constraint, in utility, at the insurance's optimal mix of $\mathbf{x} = [x_i^\circ]_i$. All firms would have to coordinate and spend more than mandated by the insurance company and hence move toward $\mathbf{x} = [x_i^\sharp]_i$ in order for the monopolist's quote to be unattractive relative to having no insurance and investing in security. Under the realistic assumption that no one firm can unilaterally change the level of attacking intensity through altruistic extra investment, the only coordinated level of security is subject to the monopolists quote $\mathbf{q} = [q_i^\circ]_i$ and is hence $\mathbf{x} = [x_i^\circ]_i$ as such the targets are a) spending more upfront on insurance than on security and b) are at a lower utility than if there was no insurance and they decided on security investments in a simultaneous Nash equilibrium.[27]

4.2. Example 2: Large diversified targets against small owner managed targets

A second example in this framework is the case when we have two types of firms in the market with differing risk preferences and expected returns to attackers. We will assume that Type 1 targets are large corporations with well diversified owners and managerial risk preferences towards risk neutral, for instance with a relative risk aversion of 1/2 or less. In contrast, Type 2 firms are smaller firms, possibly owner run, with managerial preferences tending to relatively high levels of risk aversion; for instance a coefficient of relative risk aversion between 2 and 4, see Rabin (2000) for discussion on various alternatives.

Both the number N_1 of Type 1 firms and the number N_2 of Type 2 firms are large, but N_1 is substantially smaller than N_2. The return ρ_1 on a successful attack to a Type 1 firm will be likely very high ($\rho_1 \gg 1$) whilst the maximum reward ρ_2 from a Type 2 firm may be quite low ($\rho_1 \approx 1$) and possibly even below the break even point.

In this set-up monopolist insurance is particularly unfair on Type 2 targets as their participation constraint is lower than the constraint of risk neutral Type 1 firms. Unfortunately, attackers are drawn into the market as there is a high rate of return to be gained from a successful attack on a Type 1 firm. This creates an "unsecure" environment in which Type 2 targets natural risk aversion pushes them towards (surely unfair) insurance and hence the insurer can extract proportionally more of their wealth in terms of rent.

For the quantitative derivation of the phenomenon described above, we preliminarily observe that firms within a type are otherwise identical, and act simultaneously and independently. Hence, we can simply consider the values x_k for $k \in \{1, 2\}$ as the security investment of the representative firm of each type and weight each reward ρ_k in the attacker entry condition (4) by the corresponding fraction of firms $f_k = N_k / \sum_k N_k$. The success probability σ_k has a common parameter determining its reaction to the number of attackers ($\beta_1 = \beta_2 = \beta$), as β depends essentially on the attackers' technology which this scenario assumes to be largely independent from the chosen target Type k. Still, each target Type k yields a different reward ρ_k and has a different effectiveness in security expenditures ($\alpha_1 \neq \alpha_2$).

[27] There is an alternative regulatory structure whereby a policy maker imposes a quote as a fixed proportion of L_i, for instance a fixed percentage of revenues for an online retailer, and then delegates to an insurance company the power to mandate protection and provide coverage. As the quote is fixed, the insurance company subjects target to onerous security expenditures with the attempt to drive $\sigma_i(\cdot) \to 0$ for all i, to maximize the surplus.

For notational convenience we set $\mathscr{L}_k \doteq \frac{\alpha_k-\gamma_k}{\gamma_k}(e^{\gamma_k L_k}-1)$ for $k \in \{1,2\}$ targets. Following the general treatment, we can solve the Nash equilibrium for each firm and write the optimal security investment for self protection as

$$x_k^* = \frac{1}{\alpha_k}\log \mathscr{L}_k + \frac{1}{\alpha_k}\log(n^*)^\beta \text{ where } n^* = \left(\frac{f_1\rho_1}{\mathscr{L}_1} + \frac{f_2\rho_2}{\mathscr{L}_2}\right) \tag{24}$$

It is immediate to see that even if ρ_2 is less than unity and it costs attackers more to attack them than they can earn in a reward, the attraction of $f_1\rho_1$ overwhelms this inconvenience and results in Type 2 firms having to invest significantly against randomly matched attackers attacking them.

When we introduce a monopolist insurer, the effects of any asymmetry between losses, rewards and value of the two firm types is even more pronounced. Solving the functional form in 16 for the exact quote of the monopolist for any given expenditure is given by:

$$q_k^\circ = \frac{1}{\gamma_k}\log(1 + \frac{(e^{\gamma_k L_k}-1)}{e^{\alpha_k x_k^\circ}}(n^\circ)^\beta) \text{ where } n^\circ = \left(\frac{f_1\rho_1}{e^{\alpha_1 x_1^\circ}} + \frac{f_2\rho_2}{e^{\alpha_2 x_2^\circ}}\right)^{\frac{1}{1-\beta}} \tag{25}$$

By observing that $\gamma_k < \alpha_k$, $0 \leq x_k^\circ \leq L_k$, and by taking a Taylor expansion of the quote it is possible to show that the optimal quote lies in the domain $\frac{\alpha_k-\gamma_k}{\gamma_k}L_k + \beta/\gamma_k \log(n^\circ + o(\exp(-L_k))) < q_k^\circ < L_k + \beta/\gamma_k \log(n^\circ + o(\exp(-L_k))))$ where $o(\exp(-L_k))$ is a term that exponentially decreases to zero for $L_k \to \infty$.

By construction, n° is a monotone increasing function in $f_k\rho_k$ and monotone decreasing function in $\alpha_k x_k^\circ$ we obtain to possible outcome depending on the relative expenditures and the relative value of the reward for the attackers. If the security investment x_1 of type 1 firms is sufficiently large ($e^{-\alpha_k x_k^\circ} \to 0$) and the security of small firms is negligible ($e^{-\alpha_k x_k^\circ} \to 1$), then the entire number of attackers is determined by their ability to extract small gains from a large number of small entities $n^\circ \to \mathcal{F}(f_2\rho_2)$ where $\mathcal{F}$ is a monotone function.[28] Therefore the risk fraction of the quote is essentially determined by the behavior of the small firms.

When large firms' expenditure is less effective, the reward from successful attacks to large firms $f_1\rho_1$ dominates the quote for both firm types. The implications for this scenario are quite stark: to cover their losses, due to attacks on their infrastructure that may not even be profitable ($\rho_2 < 1$), small firms will pay a disproportionate insurance premium expenditure on insurance tending to the maximal quote q_2° that they would accept.

5. Conclusions

As the problem of network security gains increasing traction in the broader policy debate, this paper provides a general treatment of the investment problem in the presence of insurance markets of differing types. The theoretical literature on insurance contracts is one of the most well developed in economics. However, cyberinsurance offers a differing set of problems to those commonly studied. Most notably, the risk generating mechanism is endogenous and driven by a strategic set of attackers. The reactivity of attackers to the investment effort of the targets is an important feature as their strategic entry and exit ensures that investment decisions by firms have a supermodular effect. Any aggregate decrease in investment results in higher ex-post attacking effort than would be expected by scaling up the marginal effect of a single firm adjustment.

[28] This behavior is consistent with the minimal conversion rate in large scale spam campaign targeting ordinary users found by Kanich et al. (2008) and the commoditization of denial of service attacks in Hutchings and Clayton (2016).

When we introduce corporate liability insurance contracts that cover losses from cyber security events targets the shift from risk averse to risk neutral decision making, almost always, results in a net reduction of investment and a net increase in attacking intensity. This effect occurs in each of the types of insurance market studied herein, unregulated but actuarially fair, regulated and actuarially fair and monopolistic. We show that there is an inherent prisoners dilemma for firms in a monopoly insurance setting. A monopolistic insurer can pitch an acceptably contract (e.g. actuarially fair insurance) and entice firms to the risk neutral point, no rational corporate officer would turn down this contract, even though the complete migration of all firms will, ceteris paribus, increase both premiums and attacks. Once ensnared, the insurance company has no incentive to increase investments to reduce attacking effort, hence nullifying the public policy objective.

We posit the most benevolent case for strategic attackers, that is they are uncoordinated and the first winner takes the entire reward. A reasonable conjecture is that refinement of the attackers capability in identifying the lottery that they are faced when mounting a campaign of attacks or differentiating their effort on the basis of the reward will normally amplify their actions[29] Another benign assumption is that attackers would leave the market under unfavorable conditions. A possible avenue for future work would be to introduce dynamic aspects and consider time-risk seekers attackers as defined in Ebert (2016): attackers who see success posted in on-line forums Allodi et al. (2016); Ooi et al. (2012) might linger in the market in the hope of future reward for a longer time than optimal given the defenders' level of investment.

Another interesting, emerging scenario is the presence of 'attacks for hire' where some hackers make available their denial of service infrastructure to on-line gamers to disrupt the play of other players, see Hutchings and Clayton (2016) for a discussion. If firms would also start to attack their competitors (e.g. when responding to on-line call for tenders) then one would need different models of 'targets', as well as radically different policy interventions. We leave this analysis for future work.

References

Allodi, L., M. Corradin, and F. Massacci (2016). Then and now: On the maturity of the cybercrime markets. *IEEE Transactions on Emerging Topics in Computing 4*(1), 35–46. Cited on p. 7, 25

Allodi, L. and F. Massacci (2014). Comparing vulnerability severity and exploits using case-control studies. *ACM Trans. Inf. Syst. Secur. 17*(1), 1:1–1:20. Cited on p. 7

Allouch, N. (2015). On the private provision of public goods on networks. *Journal of Economic Theory 157*, 527 – 552. Cited on p. 2

Baker, T. and S. J. Griffith (2007). Predicting corporate governance risk: Evidence from the directors'& officers' liability insurance market. *The university of Chicago law review*, 487–544. Cited on p. 2

Bilge, L. and T. Dumitras (2012). Before we knew it: an empirical study of zero-day attacks in the real world. In *Proc. of ACM CCS-12*, pp. 833–844. ACM. Cited on p. 7

Binmore, K. (2007). *Playing for Real*. Oxford University Press. Cited on p. 11

Bramoullé, Y., R. Kranton, and M. D'Amours (2014). Strategic interaction and networks. *American Economic Review 104*(3), 898–930. Cited on p. 2

Caillaud, B., G. Dionne, and B. Jullien (2000). Corporate insurance with optimal financial contracting. *Economic Theory 16*(1), 77–105. Cited on p. 2, 5

Cavusoglu, H., S. Raghunathan, and W. T. Yue (2008). Decision-theoretic and game-theoretic approaches to it security investment. *Journal of Management Information Systems 25*(2), 281–304. Cited on p. 6

Cornes, R. and T. Sandler (1996). *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press. Cited on p. 3

Dionne, G. and K. C. Wang (2013). Does insurance fraud in automobile theft insurance fluctuate with the business cycle? *Journal of Risk and Uncertainty 47*(1), 67–92. Cited on p. 3, 9

[29] Such cases can happen when a successful penetration of a target's network security is re-sold in a secondary market or attackers refine the group of targets that they attack using social engineering.

Ebert, S. (2016). Decision making when things are only a matter of time. Technical report, Tilburg. Available at SSRN: http://ssrn.com/abstract=2674160. Cited on p. 25

Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013, February). Selection on moral hazard in health insurance. *American Economic Review 103*(1), 178–219. Cited on p. 9

Freeman, P. and H. Kunreuther (1997). *Managing Environmental Risk Through Insurance*. Kluwer Academic Publishing. Cited on p. 3

Gordon, L. and M. Loeb (2002). The economics of information security investment. *ACM Transactions on Information and Systems Security 5*(4), 438–457. Cited on p. 6

Gordon, L. A., M. P. Loeb, and T. Sohail (2003). A framework for using insurance for cyber-risk management. *Communications of the ACM 46*(3), 81–85. Cited on p. 2

Grier, C., L. Ballard, J. Caballero, N. Chachra, C. J. Dietrich, K. Levchenko, P. Mavrommatis, D. McCoy, A. Nappa, A. Pitsillidis, et al. (2012). Manufacturing compromise: the emergence of exploit-as-a-service. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pp. 821–832. ACM. Cited on p. 7

Griffith, S. J. (2006). Uncovering a gatekeeper: Why the SEC should mandate disclosure of details concerning directors' and officers' liability insurance policies. *University of Pennsylvania Law Review*, 1147–1208. Cited on p. 2

Grossman, S. J. and O. D. Hart (1982). Corporate financial structure and managerial incentives. In *The economics of information and uncertainty*, pp. 107–140. University of Chicago Press. Cited on p. 2, 5

Herley, C. and D. Florêncio (2010). Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Economics of Information Security and Privacy*, pp. 33–53. Springer. Cited on p. 7

Hutchings, A. and R. Clayton (2016). Exploring the provision of online booter services. *Deviant Behavior*, 1–16. Cited on p. 7, 24, 25

Johnson, S. D. (2014). How do offenders choose where to offend? perspectives from animal foraging. *Legal and Criminological Psychology 19*(2), 193–210. Cited on p. 3

Kanich, C., C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage (2008). Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM conference on Computer and communications security*, pp. 3–14. Cited on p. 24

Kirwan, G. and A. Power (2013). *Cybercrime: The Psychology of Online Offenders*. Cambridge University Press. Cited on p. 3

Kotov, V. and F. Massacci (2013). Anatomy of exploit kits. In *Engineering Secure Software and Systems (ESSOS'2013)*, pp. 181–196. Springer. Cited on p. 7

Kunreuther, H. and G. Heal (2003a). Interdependent security. *The Journal of Risk and Uncertainty 26*(1), 231–249. Cited on p. 6

Kunreuther, H. and G. Heal (2003b). You only die once: Managing discrete interdependent risks. Technical report, National Bureau of Economic Research. NBER Working Paper 9885. Cited on p. 3, 6

Li, F., A. Lai, and D. Ddl (2011). Evidence of advanced persistent threat: A case study of malware for political espionage. In *Malicious and Unwanted Software (MALWARE), 2011 6th International Conference on*, pp. 102–109. Cited on p. 7

MacMinn, R. and J. Garven (2000). On corporate insurance. In *Handbook of insurance*, pp. 541–564. Springer. Cited on p. 2

Marcella Jr, A. and R. S. Greenfield (2002). *Cyber forensics: a field manual for collecting, examining, and preserving evidence of computer crimes*. CRC Press. Cited on p. 8

Mayers, D. and C. W. Smith Jr (1987). Corporate insurance and the underinvestment problem. *Journal of Risk and Insurance*, 45–54. Cited on p. 2

McCarthy, B. (2002). New economics of sociological criminology. *Annual Review of Sociology*, 417–442. Cited on p. 3

Miller, C. (2007). The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Proceedings of the 6th Workshop on Economics and Information Security*. Cited on p. 3

Moore, T. and R. Clayton (2007). Examining the impact of website take-down on phishing. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, pp. 1–13. Cited on p. 7

Moore, T. and R. Clayton (2008). The consequence of non-cooperation in the fight against phishing. In *Proceedings of the 3rd APWG eCrime Researchers Summit*. Cited on p. 7

Moore, T. and R. Clayton (2009). Evil searching: Compromise and recompromise of internet hosts for phishing. In *Proceedings of the 3rd APWG eCrime Researchers Summit*, pp. 256–272. Cited on p. 7

Nicholson, A., S. Webber, S. Dyer, T. Patel, and H. Janicke (2012). Scada security in the light of cyber-warfare. *Computers & Security 31*(4), 418–436. Cited on p. 7

O'Hearn, S. and *et al.* (2015). Insurance 2020 & beyond: Reaping the dividends of cyber resilience. Available on the web at www.pwc.com/insurance. Cited on p. 1, 5

Olakunnle, J. (2014, February). Auditing cyberinsurance policy. *ISACA Journal 2*, 29–32. Cited on p. 8

Ooi, K. W., S. H. Kim, Q.-H. Wang, and K.-L. Hui (2012). Do hackers seek variety? an empirical analysis of website defacements. In *ICIS 2012 proceedings - Research in Progress*. Cited on p. 7, 25

Pal, R., L. Golubchik, K. Psounis, and P. Hui (2013). On a way to improve cyber-insurer profits when a security vendor becomes the cyber-insurer. In *IFIP Networking Conference, 2013*, pp. 1–9. Cited on p. 4

Pauly, M. (1974). Overinsurance and public provision of insurance: The roles of moral hazard and adverse selection. *Quarterly Journal of Economics 88*(1), 44–62. Cited on p. 3

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society 32*(1/2), 122–136. Cited on p. 5

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica 68*(5), 1281–1292. Cited on p. 5, 23

Rajab, M., L. Ballard, N. Jagpal, P. Mavrommatis, N. P. D. Nojiri, and L. Schmidt (2011, July). Trends in circumventing web-malware detection. Technical report, Google. Cited on p. 7

Ransbotham, S. and S. Mitra (2009). Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research 20*. Cited on p. 6, 7

Rosen, J. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica 33*(3). Cited on p. 9

Schneier, B. (2001). *Secrets and Lies: Digital Security in a Networked World*. John Wiley and Sons. Cited on p. 2

Shavell, S. (1987). *Economic Analysis of Accident Law*. Harvard University Press. Cited on p. 3

Sood, A. K. and R. J. Enbody (2011). Malvertising–exploiting web advertising. *Computer Fraud & Security 2011*(4), 11–16. Cited on p. 7

Varian, H. (2000, 6). Managing online security risks. *New York Times*. Cited on p. 2

Verizon (2016). 2016 data breach investigation report. Technical report, Verizon. Cited on p. 7

Werlinger, R., K. Muldner, K. Hawkey, and K. Beznosov (2010). Preparation, detection, and analysis: the diagnostic work of it security incident response. *Information Management & Computer Security 18*(1), 26–42. Cited on p. 8

Electronic Companion

A. Appendix: Extended Proofs of Propositions and Theorems

We first state some helpful implications from our assumptions on how the marginal rate of the expected utility of target i changes in the number of attackers n. This marginal rate captures the changes in the perceived riskiness of the environment and is used in several of other derivations:

$$\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial n} = -\frac{\partial \sigma_i(x_i,n)L_i}{\partial n} U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i)) < 0. \tag{26}$$

The factor $U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i))$ provides the marginal utility of the target at the point of the maximally insurable loss and it is always positive and monotonically increasing in x_i as U_i is weakly convex. To prove monotonicity we consider the definition of $\mathfrak{L}$ and take the derivative of both sides. This yields the first-order derivative $\partial U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i))/\partial x_i = (-U_i'(W_i - x_i) + U_i'(W_i - x_i - L_i))/L_i$. Since U_i' is decreasing $U_i'(W_i - x_i - L_i) > U_i'(W_i - x_i)$ and therefore $\partial U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i))/\partial x_i > 0$. By assumption $\mathcal{A}.4$, $\sigma_i(x_i, n)$ increases in n and therefore the overall marginal rate is always negative for all values of x_i.

A first preliminary result we need to prove is that the variation in the number of attackers at equilibrium when the security expenditures change is negative (5). To prove it we derive both sides of the attacker entry condition in (4) and apply the global chain rule to $\partial n^\star(\mathbf{x})/\partial x_i$ by suitably aggregating σ_i and σ_j for $i \neq j$:

$$N \frac{\partial n^\star(\mathbf{x})}{\partial x_i} = \sum_{j \neq i} \rho_j \frac{\partial \sigma_j(x_j,n)}{\partial n} \frac{\partial n^\star(\mathbf{x})}{\partial x_i} + \rho_i \frac{\partial \sigma_i(x_i,n)}{\partial x_i} + \rho_i \frac{\partial \sigma_i(x_i,n)}{\partial n} \frac{\partial n^\star(\mathbf{x})}{\partial x_i}.$$

This is rearranged as follows where $n^\star(\mathbf{x})$ is replaced for n after taking the derivative by either n or x_i:

$$\frac{\partial n^\star(\mathbf{x})}{\partial x_i} = \frac{1}{\left(N - \sum_{j=1}^N \rho_j \frac{\partial \sigma_j(x_j,n)}{\partial n}\right)} \rho_i \frac{\partial \sigma_i(x_i,n)}{\partial x_i}$$

We observe that $\partial \sigma_i/\partial x_i < 0$ by Assumption $\mathcal{A}.3$. Further we can transform the denominator of the right-hand side into the partial derivative w.r.t. n of $nN - \sum_{j=1}^N \rho_j \sigma_j(x_j, n)$ where $n = n^\star(\mathbf{x})$ after taking the derivative for n. By multiplying by C, representing the nNC addendum as a sum from 1 to N of nC, and re-aggregating the summands, we obtain the negation of the second part of Assumption $\mathcal{A}.5$ on the absence of money pumps. Hence, the denominator is positive and the overall term is negative. The strict inequality of Assumption $\mathcal{A}.5$ also guarantees that $\partial n^\star(\mathbf{x})/\partial x_i$ is always well defined.

It is possible to establish a general (albeit not tight) bound on the value of the optimal investment. We consider the ratio λ_i of the marginal rate of utility in the best case scenario when no loss is present ($w = W_i$) and the unfortunate scenario where the target has spent L_i in self-protection and has been nonetheless successfully attacked ($w = W_i - 2L_i$):

$$\lambda_i \doteq \frac{U_i'(W_i)}{U_i'(W_i - 2L_i)} < 1 \text{ by Assumption } \mathcal{A}.1. \tag{27}$$

Proposition A.1
For a given number of attackers n per target, the marginal loss due to a successful attack at the equilibrium x_i^* of unregulated risk averse targets is bounded as follows:

$$-1/(\lambda_i L_i) \leq \partial(\sigma_i(x_i,n)L_i)/\partial x_i|_{x_i = x_i^*} \leq -(\lambda_i L_i). \tag{28}$$

Proof. We consider the first-order condition of the expected utility of the target in (3) and set it to zero for a given value of n_i. We obtain $\frac{\partial \sigma_i(x_i,n)}{\partial x_i} = -\mathbb{E}[U_i'(x_i)|n]/\Delta U_i(x_i)$. The concavity of $U_i(x_i)$ implies that for all $x_i \in [0, L_i]$ we have $U'(W_i - x_i) \leq \mathbb{E}[U_i'(x_i)|n] \leq U'(W_i - x_i - L_i)$ and therefore

$$-\frac{U'(W_i - x_i - L_i)}{\Delta U_i(x_i)} \leq -\frac{\mathbb{E}[U_i'(x_i)|n]}{\Delta U_i(x_i)} \leq -\frac{U'(W_i - x_i)}{\Delta U_i(x_i)}.$$

The concavity of U_i also implies that $\Delta U_i(x_i) \leq U_i'(W_i - x_i - L_i)L_i$ and therefore $(\Delta U_i(x_i))^{-1} \geq (U_i'(W_i - x_i - L_i)L_i)^{-1}$ and by multiplying both terms for the negative factor $-U_i'(W_i - x_i)$ we get $-(U_i'(W_i - x_i))/(\Delta U_i(x_i)) \leq -(U_i'(W_i - x_i))(U_i'(W_i - x_i - L_i)L_i)$. The shape of U_i also implies that $\Delta U_i(x_i) \geq U_i'(W_i - x_i)L_i$ and therefore $(U_i'(W_i - x_i)L_i)^{-1} \geq (\Delta U_i(x_i))^{-1}$. By multiplying both terms for $-U'(W_i - x_i - L_i)$ we recover $-(U'(W_i - x_i - L_i))/(U_i'(W_i - x_i)L_i) \leq -(U'(W_i - x_i - L_i))/(\Delta U_i(x_i))$. Subsequently, by joining the derived inequalities with the previous derivation we have

$$-\frac{U'(W_i-x_i-L_i)}{U_i'(W_i-x_i)L_i} \leq -\frac{U_i'(W_i-x_i-L_i)}{\Delta U_i(x_i)} \leq \quad -\frac{\mathbb{E}[U_i'(x_i)|n]}{\Delta U_i(x_i)} \quad \leq -\frac{U_i'(W_i-x_i)}{\Delta U_i(x_i)} \leq -\frac{U_i'(W_i-x_i)}{U_i'(W_i-x_i-L_i)L_i}.$$

Multiplying all terms for the positive factor L_i and replacing the expression $-\mathbb{E}[U_i'(x_i)|n]\ /\Delta U_i(x_i)$ at the center of the inequalities for the partial derivative of σ_i from (6) we obtain the desired result. □

A.1. Proof of Proposition 2.1

Proof. For the given a region of interest $[\mathbf{0}, \bar{\mathbf{x}}]$, we have to show that a solution to the first order condition $\mathbf{x}^* \in [\mathbf{0}, \bar{\mathbf{x}}]$ is a unique and stable Nash equilibrium. We derive a sufficient condition for the interval of interests given the three requirements stated in Proposition 2.1:

(i) $\mathrm{e}_{\Delta U_i} \leq -\mathrm{e}_{\partial\sigma_i/\partial n}$,

(ii) $\mathrm{e}_{\Delta U_i} \leq -1/2\mathrm{e}_{\partial\sigma_i/\partial x_i}$, and

(iii) $\mathrm{e}_{\Delta U_i(0)} \leq -\mathrm{e}_{\sigma_i(0,n_0)} - U'(W_i)/(\sigma_i(0,n_0)\Delta U_i(0))$.

The first step is proving that conditions (i–iii) imply that the marginal utility of the targets $\partial\mathbb{E}[U_i(x_i)|n]/\partial x_i$ evaluated on the constrained surface determined by the attackers' entry condition with $n = n^\star(\mathbf{x})$ is a monotone decreasing function in the interval $\mathbf{x}^* \in [\mathbf{0}, \bar{\mathbf{x}}]$.

To this extent, we represent the first-order condition for the Nash equilibrium (6) in the following form by suitably aggregating the terms of the derivative:

$$\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i} \quad = \quad -U'(W_i - x_i) - \frac{\partial\sigma_i(x_i,n)}{\partial x_i}\Delta U_i(x_i) - \sigma_i(x_i,n)\frac{\partial\Delta U_i(x_i)}{\partial x_i}$$

The partial derivative of the marginal utility can be calculated by an application of the global chain rule after having replaced $n^\star(\mathbf{x})$ for n in the right-hand side of the above equation as follows:

$$\begin{aligned}
\frac{\partial}{\partial x_i}\left(\left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}\right) = \\
U''(W_i - x_i) - \left(\frac{\partial^2\sigma_i(x_i,n)}{\partial x_i^2} + \frac{\partial^2\sigma_i(x_i,n)}{\partial x_i\partial n}\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\right)\Delta U_i(x_i) + \\
- \frac{\partial\sigma_i(x_i,n)}{\partial x_i}\frac{\partial\Delta U_i(x_i)}{\partial x_i} - \left(\frac{\partial\sigma_i(x_i,n)}{\partial x_i} + \frac{\partial\sigma_i(x_i,n)}{\partial n}\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\right)\frac{\partial\Delta U_i(x_i)}{\partial x_i} + \\
- \sigma_i(x_i,n)\frac{\partial^2\Delta U_i(x_i)}{\partial x_i^2}. \qquad (29)
\end{aligned}$$

The terms of the decomposition above can be re-arranged as follows:

$$\begin{aligned}
\ldots = &\mathbb{E}[U_i''(x_i)|n^\star(\mathbf{x})] + \\
&- \frac{\partial\sigma_i(x_i,n)}{\partial n}\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\Delta U_i(x_i)\left(\mathrm{e}_{\partial\sigma_i/\partial n} + \mathrm{e}_{\Delta U_i}\right) + \\
&- \frac{\partial\sigma_i(x_i,n)}{\partial x_i}\Delta U_i(x_i)\left(\mathrm{e}_{\partial\sigma_i/\partial x_i} + 2\mathrm{e}_{\Delta U_i}\right). \qquad (30)
\end{aligned}$$

By Assumption $\mathcal{A}.1$ we know that $\Delta U_i > 0$ for any $L_i > 0$ and $\mathbb{E}[U_i''(x_i)|n] < 0$ for any x_i and n. Consider the second addendum: Assumption $\mathcal{A}.2$ implies that $\partial\sigma_i(x_i,n)/\partial n > 0$ and therefore our preliminary result (5) implies that the addendum is not positive if $\mathrm{e}_{\partial\sigma_i/\partial n} + \mathrm{e}_{\Delta U_i(x_i)} \leq 0$, i.e.

condition (i). For the third addendum notice that Assumption $\mathcal{A}.3$ implies that $\partial\sigma_i(x_i,n)/\partial x_i < 0$ and therefore the overall addendum is negative if $\mathrm{e}_{\partial\sigma_i/\partial x_i} + 2\mathrm{e}_{\Delta U_i(x_i)} \leq 0$.

Finally, we need to show that at $\mathbf{x} = 0$ and $n = n_0$ where n_0 is the number of attackers from the Cournot subgame when the security investments of the targets are zero, $n_0 = n^\star(\mathbf{0})$, the marginal expected utility is positive, $\partial\mathbb{E}[U_i(x_i)|n]/\partial x_i|_{x_i=0,n=n_0} \geq 0$, and therefore the targets have an incentive to act and increase their expenditures.

By taking the discount factor $\sigma_i\Delta U_i$ as a common factor of the decomposition of the marginal utility we obtain the equation:

$$\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i} = -\sigma_i(x_i,n)\Delta U_i(x_i)\left(\mathrm{e}_{\sigma_i(x_i,n)} + \mathrm{e}_{\Delta U_i(x_i)} + \frac{U'(W_i-x_i)}{\sigma_i(x_i,n)\Delta U_i(x_i)}\right).$$

Observe that the product of the first two factors is always positive for all values of x_i and n. Hence the sign is only determined by the third factor. By condition (iii) it is $\mathrm{e}_{\Delta U_i(0)} \leq -\mathrm{e}_{\sigma_i(0,n_0)} - \frac{U'(W_i)}{(\sigma_i(0,n_0)\Delta U_i(0))}$. So, the third factor is also negative. Thus, $\partial\mathbb{E}[U_i(x_i)|n]/\partial x_i|_{x_i=0,n=n_0} \geq 0$. □

A.2. Proof of Proposition 2.2

Proof. We have to prove that $x_i^\dagger \geq x_i^*$.

We start by decomposing $\partial\mathbb{E}[U_P(\mathbf{x})|n^\star(\mathbf{x})]/\partial x_i$ into the two components of (9). At first we expand the definition (8)

$$\frac{\partial}{\partial x_i}\mathbb{E}[U_P(\mathbf{x})|n^\star(\mathbf{x})] = \nu_i\frac{\partial}{\partial x_i}\mathbb{E}\left[U_i(x_i)|n^\star(\mathbf{x})\right] + \textstyle\sum_{j\neq i}\nu_j\frac{\partial}{\partial x_i}\mathbb{E}\left[U_j(x_j)|n^\star(\mathbf{x})\right].$$

The next step is to expand the first addendum of the decomposition above, by replacing $\mathbb{E}\left[U_i(x_i)|n^\star(\mathbf{x})\right]$ with its definition in (3) when $n = n^\star(\mathbf{x})$. The derivative can be expanded as follows by an application of the global chain rule:

$$\begin{aligned}
\frac{\partial\mathbb{E}[U_i(x_i)|n^\star(\mathbf{x})]}{\partial x_i} &= \left(\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\left.\frac{\partial\sigma_i(x_i,n)}{\partial n}\right|_{n=n^\star(\mathbf{x})} U_i(W_i-x_i-L_i) + \right.\\
&\quad \left. + \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\left.\frac{\partial(1-\sigma_i(x_i,n))}{\partial n}\right|_{n=n^\star(\mathbf{x})} U_i(W_i-x_i)\right) +\\
&\quad + \left(\left.\frac{\partial\sigma_i(x_i,n)}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} U_i(W_i-x_i-L_i) + \sigma_i(x_i,n^\star(\mathbf{x}))\frac{\partial U_i(W_i-x_i-L_i)}{\partial x_i} + \right.\\
&\quad \left. + \left.\frac{\partial(1-\sigma_i(x_i,n))}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} U_i(W_i-x_i) + (1-\sigma_i(x_i,n^\star(\mathbf{x})))\frac{\partial U_i(W_i-x_i)}{\partial x_i}\right)\\
&= \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial n}\right|_{n=n^\star(\mathbf{x})} + \left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})}.
\end{aligned}$$

We expand the second addendum of the decomposition, again by a suitable application of the global chain rule:

$$\begin{aligned}
\sum_{j\neq i}\nu_j\frac{\partial\mathbb{E}[U_j(x_j)|n^\star(\mathbf{x})]}{\partial x_i} &= \sum_{j\neq i}\nu_j\left(\left.\frac{\partial(\sigma_j(x_j,n)U_j(W_j-x_j-L_j))}{\partial n}\right|_{n=n^\star(\mathbf{x})}\frac{\partial n^\star(\mathbf{x})}{\partial x_i} + \right.\\
&\quad \left. + \left.\frac{\partial((1-\sigma_j(x_j,n))U_j(W_j-x_j))}{\partial n}\right|_{n=n^\star(\mathbf{x})}\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\right)\\
&= \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\sum_{j\neq i}\nu_j\left.\frac{\partial\mathbb{E}[U_j(x_j)|n]}{\partial n}\right|_{n=n^\star(\mathbf{x})}.
\end{aligned}$$

By aggregating back the two addenda we obtain the desired result (9):

$$\begin{aligned}
\frac{\partial\mathbb{E}[U_P(\mathbf{x})|n^\star(\mathbf{x})]}{\partial x_i} &= \nu_i\frac{\partial n^\star(\mathbf{x})}{\partial x_i}\left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial n}\right|_{n=n^\star(\mathbf{x})} + \nu_i\left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} +\\
&\quad + \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\sum_{j\neq i}\nu_j\left.\frac{\partial\mathbb{E}[U_j(x_j)|n]}{\partial n}\right|_{n=n^\star(\mathbf{x})}\\
&= \nu_i\left.\frac{\partial\mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} + \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\left.\frac{\partial\mathbb{E}[U_P(\mathbf{x})|n]}{\partial n}\right|_{n=n^\star(\mathbf{x})}.
\end{aligned}$$

The first term of the decomposition is the value of the partial order derivative used to calculate the Nash equilibrium of the unregulated targets. In order to show that the second term is always positive we need to expand it back in a slightly different way. For brevity of exposition we omit the $n = n^\star(\mathbf{x})$ assignment:

$$\begin{aligned}\frac{\partial \mathbb{E}[U_P(\mathbf{x}|n]}{\partial n} &= \frac{\partial}{\partial n}\left(\sum_{j=1}^{N} \nu_j \mathbb{E}\left[U_j(x_j)|n\right]\right) \\ &= \sum_{j=1}^{N} \nu_j \left(\frac{\partial \sigma_j(x_j,n)}{\partial n} U_j(W_j - x_j - L_j) - \frac{\partial \sigma_j(x_j,n))}{\partial n} U_j(W_j - x_j)\right) \\ &= -\sum_{j=1}^{N} \nu_j \frac{\partial \sigma_j(x_j,n)}{\partial n} \Delta U_j(x_j).\end{aligned}$$

For every j, the component ΔU_j is always positive. Further, the marginal probability of a successful attack when the number of attackers increases $\partial \sigma_j(x_j, n)/\partial n$ is also always positive by $\mathcal{A}.2$. Hence the overall sum is always negative. This sum is multiplied by the factor $\partial n^*/\partial x_i$.

We need to further determine whether the overall second term in (9) is positive decreasing in x_i. To this extent it is useful to consider the decomposition of $\partial \mathbb{E}\left[U_P(\mathbf{x}|n^\star(\mathbf{x}))\right]/\partial n$ into the factors $-\partial n^\star(\mathbf{x})\partial x_i$ and $\sum_{j=1}^{N} \nu_j \frac{\partial}{\partial n}\sigma_j(x_j, n)\Delta U_j(x_j)$. As for the second factor we notice that $\Delta U_j(x_j)$ is positive decreasing in x_i by simple inspection and that $\partial \sigma_i(x_i, n)\partial n$ is positive by $\mathcal{A}.2$ and decreasing in x_i by $\mathcal{A}.4$. Therefore, the overall second term of the decomposition of the policy maker expected utility $\partial \mathbb{E}\left[U_P(\mathbf{x}|n^\star(\mathbf{x}))\right]/\partial x_i$ is always positive.

Then, the optimal value of $x_i^\dagger$ of the security investment of the policy maker happens at a point $(\mathbf{x}^\dagger, n^\dagger)$ where $\partial \mathbb{E}[U_i(x_i)|n]/\partial x_i|_{x_i=x_i^\dagger, n=n^\dagger} < 0$ whereas the security investment of the unregulated target happens at the place x_i^* where $\partial \mathbb{E}[U_i(x_i)|n]/\partial x_i|_{x_i=x_i^*, n=n^\dagger} = 0$. Since U_i is weakly concave then $x_i^* \leq x_i^\dagger$. □

A.3. Proof of Lemma 3.1

Proof. First, it is useful to show that risk neutral targets are indifferent to insurance. When target i is risk neutral, so that $U_{rn:i}(w) = w$, the right-hand side of (10) reduces to the quantity:

$$\begin{aligned}\mathbb{E}[U_{rn:i}(q_i, \ell_i, x_i)|\mathbf{q}, \mathbf{l}, n] &= (1 - \sigma_i(x_i, n))(W_i - x_i - q_i) + \\ &\quad \sigma_i(x_i, n)(W_i - x_i - q_i - \ell_i) \\ &= W_i - x_i - q_i - \sigma_i(x_i, n)\ell_i \\ &= W_i - x_i - \sigma_i(x_i, n)L_i.\end{aligned}$$

As in the case of no insurance discussed previously, the target's choice of defensive expenditure minimizes the expected monetary loss. Hence, $q_i = 0$ is optimal for a risk-neutral target and $L_i = \ell_i$. When target i is risk averse, (10) reduces to the quantity

$$\begin{aligned}\mathbb{E}[U_i(q_i, \ell_i, x_i)|\mathbf{q}, \mathbf{l}, n] &= (1 - \sigma_i)U_i(W_i - x_i - \sigma_i(x_i, n)(L_i - \ell_i)) + \\ &\quad \sigma_i U_i(W_i - x_i - \sigma_i(x_i, n)(L_i - \ell_i) - \ell_i) \\ &= (1 - \sigma_i)U_i(W_i - x_i - \sigma_i(x_i, n)L_i + \sigma_i\ell_i) + \\ &\quad \sigma_i U_i(W_i - x_i - \sigma_i(x_i, n)L_i + \sigma_i\ell_i - \ell_i) \\ &\leq U_i\left((1 - \sigma_i(x_i, n))(W_i - x_i - \sigma_i L_i + \sigma_i\ell_i) + \right. \\ &\quad \left. \sigma_i(W_i - x_i - \sigma_i(x_i, n)L_i + \sigma_i\ell_i - \ell_i)\right) \\ &= U_i\left(W_i - x_i - \sigma_i(x_i, n)L_i\right).\end{aligned}$$

Substituting $\ell_i = 0$ in (10), yields the result. □

A.4. Proof of Theorem 3.2

Proof. We need to prove the following statement:

$$\begin{aligned}x_i^* \geq x_i^\sharp \quad \text{iff} \quad &(i)\ \mathrm{e}_{\Delta U_i(x_i)} \leq \mathrm{e}_{\partial(1-\sigma)/\partial n} \text{ for } x_i \leq x_i^\sharp \text{ and } n = n^\star(\mathbf{x}) \\ &(ii)\ \mathbb{E}[U'(x_i^\sharp)|n^\sharp] \leq U_i'(W_i - x_i^\sharp - \mathfrak{L}(W_i - x_i^\sharp, L_i)).\end{aligned}$$

First, we consider the case in which the targets have fair insurance available. We use Proposition 3.1 to determine that targets will choose full insurance and their expected utility function is therefore

identical to $U_i(W_i - x_i - \sigma_i(x_i, n)L_i|n)$. For any given number of attackers, the maximum value of the utility will be attained by setting the usual first order condition. The derivative of the utility function, in the presence of full insurance is the following:

$$\frac{\partial U_i(W_i - x_i - \sigma_i(x_i,n)L_i)}{\partial x_i} = U_i'(W_i - x_i - \sigma_i(x_i, n)L_i)(-1 - \frac{\partial \sigma_i(x_i,n)}{\partial x_i}L_i).$$

Since U_i is positive convex, the first factor is positive for all values of wealth, i.e. $U_i' > 0$. The first-order condition can only be attained by setting the second factor to zero; this yields (7).

So we denote with $x_i^\sharp$ be the value of the security investment for the insured target, which is equal to the optimal expenditure of the risk-neutral target. For the no-insurance case the first-order condition is derived as follows:

$$\begin{aligned}
\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i} &= \frac{\partial(\sigma_i(x_i,n)U_i(W_i - x_i - L_i) + (1-\sigma_i(x_i,n))U_i(W_i - x_i))}{\partial x_i} \\
&= \frac{\partial \sigma_i(x_i,n)}{\partial x_i}U_i(W_i - x_i - L_i) + \sigma_i U_i'(W_i - x_i - L_i)(-1) + \\
&\quad -\frac{\partial \sigma_i(x_i,n)}{\partial x_i}U_i(W_i - x_i) + (1 - \sigma_i(x_i,n))U_i'(W_i - x_i)(-1) \\
&= -\sigma_i(x_i,n)U_i'(W_i - x_i - L_i) - (1 - \sigma_i(x_i,n))U_i'(W_i - x_i) + \\
&\quad -\frac{\partial \sigma_i(x_i,n)}{\partial x_i}(U_i(W_i - x_i) - U_i(W_i - x_i - L_i)) \\
&= -\mathbb{E}[U_i'(x_i)|n] - \frac{\partial \sigma_i(x_i,n)}{\partial x_i}L_i\frac{U_i(W_i - x_i) - U_i(W_i - x_i - L_i)}{L_i} \\
&= -\mathbb{E}[U_i'(x_i)|n] - \frac{\partial \sigma_i}{\partial x_i}L_i U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i)).
\end{aligned}$$

We now compute the partial derivative over n of the marginal expected utility of the risk-averse target:

$$\begin{aligned}
\frac{\partial}{\partial n}\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i} &= +\frac{\partial \sigma_i(x_i,n)}{\partial n}(U_i'(W - x_i) - U_i'(W - x_i - L_i)) - \frac{\partial}{\partial n}\frac{\partial \sigma_i}{\partial x_i}\Delta U_i \\
&= -\frac{\partial \sigma_i(x_i,n)}{\partial n}\frac{\partial \Delta U_i(x_i)}{\partial x_i} - \frac{\partial}{\partial x_i}\frac{\partial \sigma_i}{\partial n}\Delta U_i.
\end{aligned}$$

For this derivative to be greater than zero we need to impose the following condition:

$$\begin{aligned}
-\frac{\partial \sigma_i(x_i,n)}{\partial n}\frac{\partial \Delta U_i(x_i)}{\partial x_i} &> \frac{\partial}{\partial x_i}(\frac{\partial \sigma_i}{\partial n})\Delta U_i \qquad \text{moving terms on the opposite side} \\
\frac{\frac{\partial \Delta U_i}{\partial x_i}}{\Delta U_i(x_i)} &< -\frac{\frac{\partial}{\partial x_i}\partial \sigma_i(x_i,n)\partial n}{\frac{\partial \sigma_i(x_i,n)}{\partial n}} \qquad \frac{\partial \sigma_i}{\partial n} > 0 \text{ and } \Delta U_i(x_i) > 0.
\end{aligned}$$

By multiplying both terms for x_i and replacing $-\partial\sigma_i/\partial n$ with $\partial(1-\sigma_i)/\partial n$ we obtain the elasticity constraint and according to condition (i) this constraint is holding for all $x_i < x_i^\sharp$ and $n = n^\star(\mathbf{x})$. Hence the marginal expected utility in n is increasing for increasing n.

Therefore $\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i} > \frac{\partial \mathbb{E}[U_i(x_i)|n^\sharp]}{\partial x_i}$ for $n > n^\sharp$. We also know that $n > n^\sharp$ for $x_i < x_i^\sharp$ and $n = n^\star(\mathbf{x})$ By chaining the result we obtain the final inequality:

$$\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i} > \frac{\partial \mathbb{E}[U_i(x_i)|n^\sharp]}{\partial x_i} \quad \text{for } x_i < x_i^\sharp \text{ and } n = n^\star(\mathbf{x}).$$

We can now replace x_i with $x_i^\sharp$ in the right-hand side of the inequality and expand the definition of $\partial \mathbb{E}[U_i(x_i^\sharp)|n^\sharp]/\partial x_i$ and assign this value to be greater than zero. Hence we have the following inequality:

$$\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{x_i < x_i^\sharp \text{ and } n = n^\star(\mathbf{x})} > -\mathbb{E}[U_i'(x_i)|\mathbf{x}^\sharp, n^\star(\mathbf{x}^\sharp) + U_i'(W_i - x_i^\sharp - \mathfrak{L}(W_i - x_i^\sharp, L_i)) > 0.$$

Since we know that $U_i'(W_i - x_i^\sharp - \mathfrak{L}(W_i - x_i^\sharp, L_i)) > \mathbb{E}[U_i'(x_i^\sharp)|\mathbf{x}^\sharp, n^\star(\mathbf{x}^\sharp)]$ by assumption (ii) then the value of the derivative of the expected utility of the unregulated and uninsured risk-averse target is positive for all $x_i \leq *$. The value of $\partial \mathbb{E}[U_i(x_i)|n^*]/\partial x_i$ at $\mathbf{x_i}^*$ is zero for $n^* = n^\star(\mathbf{x}^*)$ by definition of optimal expenditure for the fair insurance target at equilibrium. Therefore the optimal value for which such derivative is zero will be attained at a value of x^* that is larger than $x^\sharp$. □

A.5. Proof of Proposition 3.3

Proof. First we derive equation (14) for the first-order condition:

$$\begin{aligned}
\frac{\partial \mathbb{E}[U_P(\mathbf{x})|n^\star(\mathbf{x})]}{\partial x_i} &= \sum_{j=1}^{N} \nu_i \frac{\partial U_j(W_j - x_j - \sigma_j(x_j, n^\star(\mathbf{x})L_i)}{\partial x_i} \\
&= \nu_i \frac{\partial U_i(W_j - x_j - \sigma_i(x_i, n^\star(\mathbf{x}))L_i}{\partial x_i} + \sum_{j \neq i} \nu_i \frac{\partial U_j(W_j - x_j - \sigma_j(x_j, n^\star(\mathbf{x}))L_j}{\partial x_i} \\
&= \nu_i U_i'(W_i - x_i - \sigma_i(x_i, n^\star(\mathbf{x}))) \left(-1 - \frac{\partial \sigma_i(x_i, n^\star(\mathbf{x})L_i)}{\partial x_i}\right) + \\
&\quad - \sum_{j \neq i} \nu_j U_j'(W_j - x_j - \sigma_j(x_j, n^\star(\mathbf{x}))) \frac{\partial \sigma_j(x_j, n^\star(\mathbf{x}))L_i}{\partial x_i} \\
&= \nu_i U_i'(W_i - x_i - \sigma_i(x_i, n^\star(\mathbf{x}))) \left(-1 - \frac{\partial \sigma_i(x_i, n)L_i)}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}\right. \\
&\quad \left. - \frac{\partial \sigma_i(x_i, n)L_i)}{\partial n}\Big|_{n=n^\star(\mathbf{x})} \frac{\partial n^\star(\mathbf{x})}{\partial x_i}\right) + \\
&\quad - \sum_{j \neq i} \nu_j U_j'(W_j - x_j - \sigma_j(x_j, n^\star(\mathbf{x}))) \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \frac{\partial \sigma_j(x_j, n)L_i}{\partial n}\Big|_{n=n^\star(x)} \\
&= \nu_i \frac{\partial U_i(W_i - x_i - \sigma_i(x_i, n)L_i)}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} + \\
&\quad \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \sum_{j=1}^{N} \frac{\partial U_j(W_j - x_j - \sigma_i(x_i, n))}{\partial n}\Big|_{n=n^\star(\mathbf{x})} \\
&= \nu_i \frac{\partial U_i(\mathbb{E}[x_i|n])}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} + \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \frac{\partial U_P(\mathbb{E}[\mathbf{x}|n])}{\partial n}\Big|_{n=n^\star(\mathbf{x})}.
\end{aligned}$$

Now we must establish the sign of the second term of the decomposition. The same reasoning used for proving Proposition 2.2 and Theorem 3.2. We notice that the term $\partial n^\star(\mathbf{x})/\partial x_i$ is negative as well as the term $\partial U_P(\mathbb{E}[\mathbf{x}|n])/\partial n$. Hence, their product is positive. Therefore, the value of $x_i^\sharp$ of the security investment of the policy maker happens at a point where $\partial \mathbb{E}[U_i]/\partial x_i|_{x_i=x_i^\sharp} > 0$ whereas the security investment of the unregulated target happens at the place x_i^* where $\partial \mathbb{E}[U_i]/\partial x_i|_{x_i=x_i^\sharp} = 0$. Hence $x_i^\ddagger \geq x_i^\sharp$. □ □

A.6. Proof of Theorem 3.4

Proof. First we prove that the expectation of the target is always maximized by setting deductibles $\ell_i = 0$:

$$\begin{aligned}
\mathbb{E}[U_i(q_i, \ell_i, x_i)|n] =& \sigma_i(x_i, n) U_i(W_i - x_i - q_i - \sigma_i(x_i, n)\ell_i) \\
&+ (1 - \sigma_i(x_i, n)) U_i(W_i - x_i - q_i) && (31) \\
\leq& \sigma_i(x_i, n) U_i(W_i - x_i - q_i) + (1 - \sigma_i(x_i, n)) U_i(W_i - x_i - q_i) && (32) \\
=& U_i(W_i - x_i - q_i) && (33)
\end{aligned}$$

As a next step we show that if $\mathfrak{L}(W_i - x_i, L_i) \leq q_i(\mathbf{x})$ then the marginal environmental risk $\mathfrak{R}_i(x_i, q_i, n)$ justifying additional investments for i is always smaller than the actual marginal risk $\partial \sigma_i/\partial n$ at the equilibrium point for $n = n^\star(\mathbf{x})$. At first we have $W_i - x_i - \mathfrak{L}(W_i - x_i, L_i) \geq W_i - x_i - q_i(\mathbf{x})$ and $U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i)) \leq U_i'(W_i - x_i - q)$ since $U_i' > 0$ and $U_i'' \leq 0$ by the assumption of risk aversion ($\mathcal{A}.1$). Therefore $U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i))/U_i'(W_i - x_i - q_i(\mathbf{x})) \leq 1$. Since $\partial \sigma_i(x_i, n)/\partial n \geq 0$ by Assumption ($\mathcal{A}.2$) we have that $\mathfrak{R}_i(x_i, q, n) \leq L_i \partial \sigma_i(x_i, n)/\partial n$.

Then we derive (20) and (21) from the incentive compatibility constraint (18). For the former equation we differentiate both sides of (18) and obtain

$$\frac{\partial U_i(W_i - x_i - q_i(\mathbf{x}))}{\partial x_i} = \frac{\partial \mathbb{E}[U_i(x_i)|n^\star(\mathbf{x})]}{\partial x_i}.$$

By expanding the definition of expected utility and by applying the properties of the derivation of

compound functions we get

$$\begin{aligned} U_i'(W_i - x_i - q_i)(-\partial q_i/\partial x_i - 1) = & \frac{\partial \sigma_i(x_i, n^\star(\mathbf{x}))}{\partial x_i} U_i(W_i - x_i - L_i)) \\ & + \sigma_i(x_i, n^\star(\mathbf{x})) \frac{\partial U_i(W_i - x_i - L_i)}{\partial x_i} + \\ & + \frac{\partial (1 - \sigma_i(x_i, n^\star(\mathbf{x}))}{\partial x_i} U_i(W_i - x_i) \\ & + (1 - \sigma_i(x_i, n^\star(\mathbf{x}))) \frac{\partial U_i'(W_i - x_i)}{\partial x_i}. \end{aligned} \tag{34}$$

We can now process the right-hand side term of the equation by using the global chain rule and by re-arranging terms:

$$\begin{aligned} \dots &= \left(\tfrac{\partial \sigma_i(x_i,n)}{\partial x_i} + \tfrac{\partial \sigma_i(x_i,n)}{\partial n} \tfrac{\partial n^\star(\mathbf{x})}{\partial x_i} \right)_{n=n^\star(\mathbf{x})} \cdot U_i(W_i - x_i - L_i)) + \\ &\quad + \left(\tfrac{\partial (1-\sigma_i(x_i,n))}{\partial x_i} + \tfrac{\partial (1-\sigma_i(x_i,n))}{\partial n} \tfrac{\partial n^\star(\mathbf{x})}{\partial x_i} \right)_{n=n^\star(\mathbf{x})} \cdot U_i(W_i - x_i) \\ &\quad + \sigma_i(x_i, n^\star(\mathbf{x})) \tfrac{\partial U_i(W_i - x_i - L_i)}{\partial x_i} + (1 - \sigma_i(x_i, n^\star(\mathbf{x})) \tfrac{\partial U_i(W_i - x_i)}{\partial x_i} \\ &= \tfrac{\partial \sigma_i(x_i,n)}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} U_i(W_i - x_i - L_i)) + \sigma_i(x_i, n^\star(\mathbf{x})) \tfrac{\partial U_i(W_i - x_i - L_i)}{\partial x_i} + \\ &\quad + \tfrac{\partial (1-\sigma_i(x_i,n))}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} U_i(W_i - x_i) + (1 - \sigma_i(x_i, n^\star(\mathbf{x})) \tfrac{\partial U_i(W_i - x_i)}{\partial x_i} \\ &\quad + \tfrac{\partial n^\star(\mathbf{x})}{\partial x_i} \tfrac{\partial \sigma_i(x_i,n)}{\partial n}\Big|_{n=n^\star(\mathbf{x})} U_i(W_i - x_i - L_i)) + \\ &\quad + \tfrac{\partial n^\star(\mathbf{x})}{\partial x_i} \tfrac{\partial (1-\sigma_i(x_i,n))}{\partial n}\Big|_{n=n^\star(\mathbf{x})} U_i(W_i - x_i). \end{aligned}$$

The first four terms can be aggregated back into the partial derivative of the expected utility according to (3) where n is held constant during the derivative and then replaced by $n^\star(\mathbf{x})$. The remaining two terms can be aggregated into the definition of $\Delta U_i(x_i)$ after changing the sign accounting for the negative sign stemming from $\partial(1 - \sigma_i)/\partial n$:

$$\begin{aligned} U_i'(W_i - x_i - q_i(\mathbf{x})) \left(-1 - \frac{\partial q_i(\mathbf{x}))}{\partial x_i} \right) = & \frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} + \\ & - \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \frac{\partial \sigma_i(x_i, n)}{\partial n}\Big|_{n=n^\star(\mathbf{x})} \Delta U_i(x_i). \end{aligned} \tag{35}$$

By multiplying and dividing the right-hand side by L_i, and by replacing the definition of maximally insurable loss from equation (2) in terms of $\Delta U_i(x_i)/L_i$ one obtains the following equation after simplification of the left-hand side:

$$\begin{aligned} \left(-1 - \frac{\partial q_i(\mathbf{x}))}{\partial x_i} \right) = & \frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} \cdot \frac{1}{U_i'(W_i - x_i - q_i(\mathbf{x}))} + \\ & - \frac{\partial n^\star(\mathbf{x})}{\partial x_i} \cdot \frac{\partial \sigma_i(x_i, n)}{\partial n}\Big|_{n=n^\star(\mathbf{x})} \cdot L_i \cdot \frac{U_i'(W_i - x_i - \mathfrak{L}(W_i - x_i, L_i))}{U_i'(W_i - x_i - q_i(\mathbf{x}))}. \end{aligned} \tag{36}$$

Replacing the definition of marginal environment risk yields our desired decomposition as given in (20).

For the second equation (21) we use the following derivation:

$$\begin{aligned}
\tfrac{\partial U_j(W_j-x_j-q_j(\mathbf{x}))}{\partial x_i} &= \tfrac{\partial}{\partial x_i}\mathbb{E}[U_j(x_j)|n^\star(\mathbf{x})] \\
U_j'(W_j-x_j-q_j(\mathbf{x}))\tfrac{\partial (W_j-x_j-q_j(\mathbf{x}))}{\partial x_i} &= \tfrac{\partial}{\partial x_j}(\sigma_j(x_j,n^\star(\mathbf{x}))U_j(W_j-x_j-L_j)) \\
&\quad +(1-\sigma_j(x_j,n^\star(\mathbf{x})))U_j(W_j-x_j) \\
U_j'(W_j-x_j-q_j(\mathbf{x}))\left(-\tfrac{\partial q_j(\mathbf{x})}{\partial x_i}\right) &= \tfrac{\partial \sigma_j(x_j,n^\star(\mathbf{x}))}{\partial x_i}U_j(W_j-x_j-L_j))+ \\
&\quad +\tfrac{\partial (1-\sigma_j(x_j,n^\star(\mathbf{x}))}{\partial x_i}U_j(W_j-x_j) \\
&= \tfrac{\partial \sigma_j(x_j,n^\star(\mathbf{x}))}{\partial x_i}\left(U_j(W_j-x_j-L_j))-U_j(W_j-x_j)\right) \\
&= \tfrac{\partial \sigma_j(x_j,n)}{\partial n}\Big|_{n=n^\star(\mathbf{x})} \\
&\quad \cdot\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\left(U_j(W_j-x_j-L_j))-U_j(W_j-x_j)\right).
\end{aligned}$$

Finally, by multiplying and dividing by L_j and reverting to the definition of maximally insurable loss we obtain the final result:

$$\begin{aligned}
U_j'(W_j-x_j-q_j(\mathbf{x}))\tfrac{\partial q_j(\mathbf{x})}{\partial x_i} &= \tfrac{\partial \sigma_j(x_j,n)L_j}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}\cdot\tfrac{U_j(W_j-x_j)-U_j(W_j-x_j-L_j))}{L_j} \\
\tfrac{\partial q_j(\mathbf{x})}{\partial x_i} &= \tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\cdot\tfrac{\partial \sigma_j(x_j,n)L_j}{\partial n}\Big|_{n=n^\star(\mathbf{x})}\cdot\tfrac{U_j'(W_j-x_j-\mathfrak{L}(W_j-x_j))}{U_j'(W_j-x_j-q_j(\mathbf{x}))}.
\end{aligned}$$

Now the insurer will optimize its profit function by taking the usual first-order condition:

$$\begin{aligned}
\tfrac{\partial \Pi(\mathbf{x})}{\partial x_i} &= \tfrac{\partial \sum_{j=1}^N q_j(\mathbf{x})-\sigma_j(x_j,n^\star(\mathbf{x})))L_j}{\partial x_i} \\
&= \tfrac{\partial (q_i(\mathbf{x})-\sigma_i(x_i,n^\star(\mathbf{x}))L_i)}{\partial x_i}+\textstyle\sum_{j\neq i}\tfrac{\partial q_j(\mathbf{x})-\sigma_j(x_j,n^\star(\mathbf{x}))L_j}{\partial x_i}.
\end{aligned}$$

We expand the first term of the decomposition

$$\begin{aligned}
\tfrac{\partial q_i(\mathbf{x})-\sigma_i(x_i,n^\star(\mathbf{x}))L_i)}{\partial x_i} &= \tfrac{\partial q_i(\mathbf{x})}{\partial x_i}-\tfrac{\partial \sigma_i(x_i,n^\star(\mathbf{x}))L_i)}{\partial x_i} \\
&= \tfrac{\partial q_i(\mathbf{x})}{\partial x_i}-\left(\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial n}\Big|_{n=n^\star(\mathbf{x})}\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\right) \\
&= -1-\tfrac{1}{U_i'(W_i-x_i-q_i(x_i))}\tfrac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\mathfrak{R}_i(x_i,q_i,n^\star(\mathbf{x})) \\
&\quad -\left(\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial n}\Big|_{n=n^\star(\mathbf{x})}\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\right) \\
&= -\tfrac{1}{U_i'(W_i-x_i-q_i(x_i))}\tfrac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+ \\
&\quad -1-\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+ \\
&\quad -\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\left(\tfrac{\partial \sigma_i(x_i,n)L_i}{\partial n}\Big|_{n=n^\star(\mathbf{x})}-\mathfrak{R}_i(x_i,q_i,n^\star(\mathbf{x}))\right),
\end{aligned}$$

expanding the second term for every addendum of the sum

$$\begin{aligned}
\tfrac{\partial q_j(\mathbf{x})-\sigma(x_j,n^\star(\mathbf{x}))L_j}{\partial x_i} &= \tfrac{\partial q_j(\mathbf{x})}{\partial x_i}-\tfrac{\partial \sigma(x_j,n^\star(\mathbf{x}))L_j}{\partial x_i} \\
&= \tfrac{\partial q_j(\mathbf{x})}{\partial x_i}-\tfrac{n^\star(\mathbf{x})}{\partial x_i}\tfrac{\sigma_j(x_j,n)L_j}{n}\Big|_{n=n^\star(\mathbf{x})} \\
&= -\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\cdot\left(\tfrac{\partial \sigma_j(x_j,n)L_j}{\partial n}\Big|_{n=n^\star(\mathbf{x})}-\mathfrak{R}_j(x_j,q_j,n^\star(\mathbf{x}))\right).
\end{aligned}$$

We can now start to group terms appropriately:

$$\begin{aligned}
\tfrac{\partial \Pi}{\partial x_i} &= -\tfrac{1}{U_i'(W_i-x_i-q_i(x_i))}\tfrac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})}+ \\
&\quad -1-\tfrac{\partial \sigma(x_i,n)L_i}{\partial x_i}\Big|_{n=n^\star(\mathbf{x})} \\
&\quad -\tfrac{\partial n^\star(\mathbf{x})}{\partial x_i}\cdot\textstyle\sum_{j=1}^N\left(\tfrac{\partial \sigma_j(x_j,n)L_j}{\partial n}\Big|_{n=n^\star(\mathbf{x})}-\mathfrak{R}_j(x_j,q_j,n^\star(\mathbf{x}))\right).
\end{aligned}$$

If we replace the definition of enviromental risk we obtain the desired result:

$$\begin{aligned}\frac{\partial \Pi}{\partial x_i} &= -\frac{1}{U_i'(W_i - x_i - q_i(\mathbf{x}))} \left.\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} + \\ & -1 - \left.\frac{\partial \sigma(x_i,n)L_i}{\partial x_i}\right|_{n=n^\star(\mathbf{x})} \\ & -\frac{\partial n^\star(\mathbf{x})}{\partial x_i} \cdot \sum_{j=1}^{N} \left.\frac{\partial \sigma_j(x_j,n)L_j}{\partial n}\right|_{n=n^\star(\mathbf{x})} \frac{U_j'(W_j - x_j - \mathfrak{L}(W_j - x_j))}{U_j'(W_j - x_j - q_j(\mathbf{x}))}\end{aligned}$$

We consider now the value of the above derivative for $x_i = x_i^\circ$, the optimal value for the insurer. Since the profit of the insurer reaches the maximum the left hand side must be equal to zero and we can move the marginal expected value of the target for the Nash equilibrium on the left-hand side of the equation:

$$\begin{aligned}\frac{1}{U_i'(W_i - x_i^\circ - q_i(\mathbf{x}^\circ))} \left.\frac{\partial \mathbb{E}[U_i(x_i)|n]}{\partial x_i}\right|_{x_i=x_i^\circ, \mathbf{x}=\mathbf{x}^\circ, n=n^\star(\mathbf{x}^\circ)} &= -1 + \\ - \left.\frac{\partial \sigma(x_i,n)L_i}{\partial x_i}\right|_{x_i=x_i^\circ, n=n^\star(\mathbf{x}^\circ)} & \qquad (37) \\ - \frac{\partial n^\star(\mathbf{x}^\circ)}{\partial x_i} \cdot \sum_{j=1}^{N} \left.\frac{\partial \sigma_j(x_j,n)L_j}{\partial n}\right|_{x_j=x_j^\circ, n=n^\star(\mathbf{x}^\circ)} & \qquad (38) \\ \cdot \frac{U_j'(W_j - x_j^\circ - \mathfrak{L}(W_j - x_j^\circ))}{U_j'(W_j - x_j^\circ - q_j(\mathbf{x}^\circ))} & \qquad (39)\end{aligned}$$

We apply the same reasoning used for the proof of Proposition 2.1 and Theorem 3.2 to derive the inequality for the value of marginal expected utility at the Nash equilibrium. If all elasticity conditions for Proposition 2.1 hold for $x_i \leq x_i^\circ$ then we know that the marginal expected utility of the unregulated target is monotone decreasing for all values $x_i \leq x_i^\circ$. If the marginal expected loss at x_i° is smaller than -1, that is $\partial \sigma(x_i,n)/\partial x_i|_{x_i=x_i^\circ, n=n^\star(\mathbf{x}^\circ)} \leq -1/L_i$, then the first line of the equation is a positive term which is also monotone decreasing. The third term is always positive given assumption $\mathcal{A}.1$ and therefore $x_i^\circ \leq x_i^*$. □

B. Appendix: Example Calculations

B.1. Derivation for homogeneous firm

Plugging the functional forms for $U(\cdot)$ and $\sigma(\cdot)$ into Propositions 2.1, Theorem 3.2 and Theorem 3.4 and rearranging we derive the solution to the expected attackers per target: $n^\star(\mathbf{x}) = (\exp(-\alpha x)\rho)^{1/(1-\beta)}$; the Nash equilibrium investment in the absence of insurance: $x^* = \frac{\beta}{\alpha}\log(\rho) + \frac{1-\beta}{\alpha}\log(\alpha L - \gamma L) + \frac{1-\beta}{\alpha}\cdot\gamma\cdot\mathfrak{L}(w,L)$; the investment under fair insurance/risk neutral investment: $x^\sharp = \frac{\beta}{\alpha}\log(\rho) + \frac{1-\beta}{\alpha}\log(\alpha L)$; the monopolist insurer optimal quote: $q^\circ = \frac{1}{\gamma}\log(\exp(\gamma L)n^\beta - n^\beta + \exp(\alpha x)) - \frac{\alpha x}{\gamma}$ and the monopolist insurers desired level of investment: $x^\circ = x^\sharp - \frac{1-\beta}{\alpha}(\log(\frac{\alpha}{\gamma}L) + \log(1-\exp(-\gamma\mathfrak{L}(W-x^\circ,L)))) = x^* - \frac{1-\beta}{\alpha}(\log(\frac{\alpha}{\gamma}L - L) + \gamma\mathfrak{L}(w,L) + \log(1-\exp(-\gamma\mathfrak{L}(W-x^\circ,L))))$.

C. Appendix: Actual Insurance Contracts and Estimation of Risks

Cyberinsurance has only recently been made available widely. It is therefore useful to gain some insight into the likelihoods of losses and their magnitude.

Table 1 provides a series of quotes from a major insurer for a range of different firms. Cross-sectionally the size of the coverage tallies with the distribution of claims reported in a survey of companies by the Net Diligence organization,[30] which surveys 16 insurance companies annually regarding cyberinsurance claims. The group of insurers reported 85 claims for US firms that resulted in payouts for the 2011 to 2013 period with claims ranging from $1,000 to $13.7 million, with a median of around $750,000. The claim survey information is consistent with the indicative quotes in that the approximate size of payouts in most sectors was just under 1 million dollars and this is the approximate level of coverage. However, certain vulnerable sectors choose a coverage well in excess of this amount and this is borne out by the realized claims that included several that were in excess of $10 million.

Using these quotes we have made some illustrative calculations to estimate the natural probability of a successful attack. First we divide the quoted premium by the coverage limit to derive the probability of an incident, in one year, under the assumption that the quote is actuarially fair. This corresponds to column 5 in Table 1. We can see that the highest probability under this assumption is for financial and E-commerce firms at 3.7%.

This is likely a misleading value as the insurance market in this area is, in most likelihood, far away from being actuarially fair. We therefore further assume that the insurance company, as a near monopolist, can charge a monopoly price up to the break-even of expected utility for the firm versus the certain utility in presence of insurance. For illustration purposes, we presume that security expenditure is the same for both insured and uninsured targets and assume that the target preferences are described by a CARA utility function as in Section 4: $\tilde{U}_i(z) = -(1/\tilde{\gamma})e^{-\tilde{\gamma}z}$, where $\tilde{\gamma} = 2$. We then compute the value of $\tilde{\sigma}_i$ such that targets are indifferent between taking insurance and staying uninsured: $\tilde{U}_i(1-q_i/W_i) = \sigma_i U_i(1-L_i/W_i) + (1-\sigma_i)U_i(1)$. The corresponding value is reported in column 6 in Table 1.

We also evaluate the procedure for a variety of risk aversion coefficients from 0.1 to 1 while considering losses are relative to annual total revenue, which is a standard approach for for corporate liability insurance. We choose 0.1, as a control value, as the probability should be very close to the actuarially fair insurance. This is indeed the case for all organizations in this sample. For a risk aversion coefficient of 1, the probability of an event with a successful claim for the Financial and E-commerce sector is essentially identical to the 0.1 case. However, for several other firms, the implied probability of a claim event rise by a five-fold factor.

In a final experiment we used two industry-reported payout ratios of 10% and 50% on premiums and compute the minimum implied constant relative risk aversion for firms with iso-elastic power utility, by reversing the calculation used above.[31] For the widely quoted 10% payout, all firms have a risk aversion coefficient well above unity; however, for a more reasonable 50% payout several firms have minimum relative risk aversion coefficients close to a half.

[30] Article *Net Diligence Cyber Claims Study* 2014 by Mark Greisiger.

[31] See: "Cyber insurance market tempts new participants" by Alistair Grey, *Financial Times*, October 6, 2014. `http://www.ft.com/cms/s/0/69db580c-4d37-11e4-8f75-00144feab7de.html`.

Table 1: Selection of cyberinsurance contracts across a variety of commercial settings.

Industry	$\sim W_i$	L_i	q_i	q/L	$\tilde{\sigma}_i$
Healthcare	25,000,000	1,000,000	12,900	1.29%	1.24%
Education	25,000,000	1,000,000	6,000	0.60%	0.58%
Financial	100,000,000	1,000,000	37,000	3.70%	3.66%
Retail	50,000,000	1,000,000	26,000	2.60%	2.55%
E-commerce	50,000,000	1,000,000	37,000	3.70%	3.63%
Restaurant Chain	50,000,000	1,000,000	10,000	1.00%	0.98%
Manufacturing	100,000,000	10,000,000	50,000	0.50%	0.45%
Healthcare IT	1,200,000	5,000,000	15,900	0.32%	0.00%
Healthcare SaaS	1,500,000	5,000,000	30,420	0.61%	0.01%
Electronic Health Records Stor.	5,000,000	1,000,000	8,010	0.80%	0.65%
Clinical Data	20,000	2,000,000	4,900	0.25%	0.00%
E-Waste Company	1,500,000	1,000,000	3,564	0.36%	0.17%
Psychologists Office	1,000,000	1,000,000	1,600	0.16%	0.05%
Doctor's Office	700,000	500,000	649	0.13%	0.06%
Online Retailer	500,000	1,000,000	1,100	0.11%	0.01%
Hospital	170,000,000	5,000,000	42,000	0.84%	0.82%
Data Storage	15,000,000	20,000,000	120,000	0.60%	0.12%

Notes: The first column identifies the industrial sector for which the corporate liability insurance has a specific cyberinsurance clause. The second column denotes the reported revenue (approximated by the insurance company) of the organization in one year proxying for W_i. The third column provides the level of coverage L_i for each organization while the fourth column reports the payable insurance premium q_i. From this data we computed column five as the ratio of the premium to the coverage: if the insurance was actuarially fair, then this would be the probability of a successful attack, assuming one claim per year. Column six is the probability that is estimated from the data by assuming that targets are indifferent to insurance and have a CARA utility function with a risk aversion coefficient equal to 2.