

Cyber Insurance as an Incentive for Internet Security

Jean Bolot
Sprint, California, USA
bolot@sprint.com

Marc Lelarge
INRIA-ENS, Paris, France
marc.lelarge@ens.fr

Abstract

Managing security risks in the Internet has so far mostly involved methods to reduce the risks and the severity of the damages. Those methods (such as firewalls, intrusion detection and prevention, etc) reduce but do not eliminate risk, and the question remains on how to handle the residual risk. In this paper, we consider the problem of whether buying insurance to protect the Internet and its users from security risks makes sense, and if so, of identifying specific benefits of insurance and designing appropriate insurance policies.

Using insurance in the Internet raises several questions because entities in the Internet face correlated risks, which means that insurance claims will likely be correlated, making those entities less attractive to insurance companies. Furthermore, risks are interdependent, meaning that the decision by an entity to invest in security and self-protect affects the risk faced by others. We analyze the impact of these externalities on the security investments of the users using simple models that combine recent ideas from risk theory and network modeling.

Our key result is that using insurance would increase the security in the Internet. Specifically, we show that the adoption of security investments follows a threshold or tipping point dynamics, and that insurance is a powerful incentive mechanism which pushes entities over the threshold into a desirable state where they invest in self-protection.

Given its many benefits, we argue that insurance should become an important component of risk management in the Internet, and discuss its impact on Internet mechanisms and architecture.

presented at: WEIS 2008¹, Seventh Workshop on the Economics of Information Security, Hanover NH (USA), June 25-28, 2008.

¹shortened version presented at INFOCOM 08 (mini-Conference) [5].

1 Introduction

The Internet has become a strategic infrastructure in modern life and as such, it has become critical to the various entities (operators, enterprises, individuals,...) which deliver or use Internet services to protect that infrastructure against risks. The four typical options available in the face of risks are to: 1) avoid the risk, 2) retain the risk, 3) self-protect and mitigate the risk, and 4) transfer the risk. Option 1 involves preventing any action that could involve risk, and it is clearly not realistic for the Internet. Option 2 involves accepting the loss when it occurs. Option 3 involves investing in methods to reduce the impact of the risk and the severity of the damages. Option 4 involves transferring the risk to another willing party through contract or hedging.

Most entities in the Internet have so far chosen, or are only aware of the possibility of, a mix of options 2 and 3. As a result, these entities have been busy investing in people and devices to identify threats and develop and deploy countermeasures. In practice, this has led to the development and deployment of a vast array of systems to detect threats and anomalies (both malicious such as intrusions, denial-of-service attacks, port scanners, worms, viruses, etc., and non-intentional such as overloads from flash crowds) and to protect the network infrastructure and its users from the negative impact of those anomalies, along with efforts in the area of security education in an attempt to minimize the risks related to the human factor [10]. In parallel, most of the research on Internet security has similarly focused on issues related to option 3, with an emphasis on algorithms and solutions for threat or anomaly detection, identification, and mitigation.

However, **self protecting against risk or mitigating risk does not eliminate risk**. There are several reasons for this. First, there do not always exist fool-proof ways to detect and identify even well defined threats; for example, even state of the art detectors of port scanners and other known anomalies suffer from non-zero rates of false positives and false negatives [30]. Furthermore, the originators of threats, and the threats they produce, evolve on their own and in response to detection and mitigation solutions being deployed, which makes it harder to detect and mitigate evolving threat signatures and characteristics [54]. Other types of damages caused by non-intentional users, such as denial of service as a result of flash crowds, can be predicted and alleviated to some extent but not eliminated altogether. Finally, eliminating risks would require the use of formal methods to design provably secure systems, and formal methods capture with difficulty the presence of those messy humans, even non malicious humans, in the loop [45].

In the end, despite all the research, time, effort, and investment spent in Internet security, there remains a residual risk: the Internet infrastructure and its users are still very much at risk, with accounted damages already reaching considerable amounts of money and possible damage even more daunting (e.g. [24], [55] for a discussion on worm damage and conference web site for an opinion on damage cost estimation.) **The question then is how to handle this residual risk.**

One way to handle residual risk which has not been considered in much detail yet is to use the fourth option mentioned above, namely transfer the risk to another willing entity through contract or hedging. A widely used way to do this is through insurance, which is one type of risk transfer using contracts. In practice, the risk is transferred to an insurance company, in return for a fee which is the insurance premium. Insurance allows individuals or organizations to smooth payouts for uncertain events (variable costs of the damages associated with security risks) into predictable periodic costs. **Using insurance to handle security risks in the Internet raises several questions: does this option make sense for the Internet, under which circumstances? Does it provide benefits, and if so, to whom, and to what extent?** Our goal in this paper is to consider those questions.

There have traditionally been two approaches to modeling insurance and computing premiums, an actuarial approach and an economic approach. The actuarial approach uses the classical model for insurance risk where the risk process $U(t)$ is expressed as

$$U(t) = C + \varphi t - S(t), \quad t \geq 0, \quad (1)$$


where C is the initial capital, φ is the premium rate and the claim amount $S(t) = \sum_i X_i$ consists of a random sum of claims X_i , $1 \leq i \leq N(t)$ where $N(t)$ is the number of claims until time t . The goal of the modeling effort is, given statistics on the claims, to determine a premium rate φ which avoids the so-called ruin for the insurer, i.e. a negative value of $U(t)$ (for a large initial capital C). Simple models consider for $\{N(t)\}$ a homogeneous Poisson process. To capture the correlation between risks

faced by users, and therefore between claims made by those users, some approaches model claims using heavy-tailed distributions (refer to the textbook [42] for details). In [26], Herath et al. use an actuarial approach to price the premium based on copula methodology.

The economic approach considers that a limit to insurability cannot be defined only on the characteristics of the risk distribution, but should take into account the economic environment. We take this approach in the paper. We consider a sequence of increasingly complex, but simple models, to examine the impact of insurance in the Internet.

Our first model is the classical, expected utility model with a single entity or user. We use it to present known results from the literature, and in particular to examine the interplay between self-protection and insurance. The main relevant result is that the insurance premium should be negatively related to the amount invested by the user in security (self-protection). This parallels the real life situation where homeowners who invest in a burglar alarm and new locks expect their house theft premium to decrease following their investment.

The single user model is not appropriate for our purpose because the entities in the Internet face risks that are *correlated*, meaning that the risk faced by an entity increases with the risk faced by the entity's neighbors (e.g. I am likely to be attacked by a virus if my neighbors have just been attacked by that virus). Furthermore, entities face risks that are *interdependent*, meaning that those risks depend on the behavior of other entities in the network (such as their decisions to invest in security). Thus, the reward for a user investing in security depends on the general level of security in the network, leading to the feedback loop situation shown below.



We analyze the impact of these externalities on the security investments of the users with and without insurance being available. We focus on risks such as those caused by propagating worms or viruses, where damages can be caused either directly by a user, or indirectly via the user's neighbors. Users can decide whether or not to invest some amount c in security solutions to protect themselves against risk, which eliminates direct (but not indirect) damages. In the 2-user case, Kunreuther and Heal [36] proved that, in the absence of insurance, there exists a Nash equilibrium in a "good" state (where both users self protect) if the security investment cost c is low enough. These results were recently extended by the authors to a network setting in [37] and [38].

We first build upon this result to add insurance to the 2-user case. We then consider the general case of a n -user network for which damages spread among the users that decide whether or not to invest in security for self-protection. We compare both situations when insurance is available and when it is not. We show that **if the premium discriminates against user that do not invest in security, then insurance is a strong incentive to invest in security.** We also show how **insurance can be a mechanism to facilitate the deployment of security investments by taking advantage network effects such as threshold or tipping point dynamics.**

The models we use in the paper are simple, and our results will not by themselves establish insurance markets for the Internet and its users. Still, the models and results are significant because they provide a convenient way to formulate the problem of deploying insurance in the Internet, they provide a methodology to evaluate the impact of insurance and design appropriate insurance policies, and they bring out the significant benefits of insurance. Given those benefits, we believe that insurance ought to be considered as an important component of Internet security, as a mechanism to increase the adoptability of security measures Internet-wide, and as a mechanism that could have significant impact on Internet architecture and policies.

The rest of the paper is organized as follows. In Section 2, we describe related work. In Section 3, we introduce the classical expected utility model for a single user and present the standard results about risk premium and the interplay between self-protection and insurance. In Section 4, we describe the 2-user model, present the known results for self-protection in the absence of insurance, then build on this model to include insurance and prove our main results in the 2-user case. In Section 5, we extend those results to the case of a general network of n users. In Section 6, we discuss the impact of insurance and risk transfer on Internet mechanisms and architecture. Section 7 concludes the paper.

2 Related work

Risk management in the Internet has typically involved approaches that retain the risk (i.e. accept the loss when it occurs) and self-protect against the risk. As a result, a vast amount of research has been published in the area of protection against risk in the Internet, ranging from risk or threat detection, identification, mitigation, to ways to survive or recover from damages (refer to the large body of research published in related conferences [29], and in relevant security conferences [28, 43]). In parallel, researchers in the insurance community published a vast body of results in the area of insurance against risk (e.g. [21, 18]).

Comparatively little has been carried out or published at the intersection of insurance and the Internet. We can divide relevant contributions in three areas: **Internet economics** (without insurance), **cyberinsurance** or insurance of computer risks in general (without much focus on network effects), and **insurance of correlated or interdependent risks**.

Research on **Internet economics** aims at increasing our understanding of the Internet as an economic system and at developing policies and mechanisms to achieve desirable economic goals (much the same way early research on the Internet aimed at developing policies and mechanisms - such as the IP protocol - to achieve desirable design goals such as those described in [12], or more recent research aims at developing clean-slate policies and mechanisms to achieve the desired goals of the future Internet [19]). The importance of the economic aspects of the Internet was recognized very early on. Kleinrock in 1974 mentioned that "[H]ow does one introduce an equitable charging and accounting scheme in such a mixed network system. In fact, the general question of accounting, privacy, security and resource control and allocation are really unsolved questions which require a sophisticated set of tools" [35]. More recently, Clark et al [13] mention economic drivers as key drivers to revisit old design principles and suggest new ones. Research in Internet economics has examined several issues, such as the economics of digital networks (refer to [53] for pointers to recent work in the area, and e.g. [22] for the analysis of a point problem, specifically the impact of layering), pricing models and incentive mechanisms for resource allocation that align the interests of possibly selfish users with the interests of the network architect [52, 40, 31], and the economics of security (refer to [2] for a recent survey and references, also [8] and the proceedings of the Workshop on economics of information security).

Using **cyberinsurance** as a way to handle the residual risk after computer security investments have been made was proposed more than 10 years ago in the computer science literature [39] but popularized only recently by Schneier [50, 51]. The problem of residual risk and cyber insurance has been analyzed by Gordon et al. in [25]. Kesan et al. in [33, 34] make the economic case for insurance, arguing that insurance results in higher security investments (and therefore increases the global level of safety), that it encourages standards for best practices to be at the socially optimum level, and that it solves a market failure (namely the absence of risk transfer opportunity), and they see the emerging market for cyberinsurance as a validation of the case they make in the paper.

The market for cyberinsurance started in the late 90's with insurance policies offered by security software companies partnering with insurance companies as packages (software + insurance). The insurance provided a way to highlight the (supposedly high) quality of the security software being sold, and to deliver a "total" risk management solution (risk reduction + residual risk transfer), rather than the customary risk reduction-only solution (combined with risk retaining); see for examples solutions offered by Cigna (Cigna's Secure System Insurance) or Counterpane/Lloyd's of London [15]. More recently, insurance companies started offering standalone products (e.g. AIG's NetAdvantage [1]). Majuca et al. [41] provide a recent and comprehensive description of the history and the current state of computer insurance.

A challenging problem for Internet insurance companies is caused by **correlations between risks**, which makes it difficult to spread the risk across customers - a sizeable fraction of worm and virus attacks, for example, tend to propagate rapidly throughout the Internet and inflict correlated damages to customers worldwide [56, 48]. Furthermore, entities in the Internet face **interdependent risks**, i.e. risks that depend on the behavior of other entities in the network (e.g. whether or not they invested in security solutions to handle their risk), and thus the reward for a user investing in security depends on the general level of security in the network. Correlated and interdependent risks have only very recently started being addressed in the literature. Böhme in [6] considers insurance with correlations in the extreme case of a monoculture (a system of uniform agents) with correlated Bernoulli risks and argues that the

strong correlation of claims in that case may indeed hinder the development of a cyberinsurance industry. Subsequent work in [7] argues that correlations are actually two-tiered and supports the argument with honeypot data. One tier represents the correlations across risks within an entity such as a corporation, the other tier represents the correlations of risks across independent entities. Correlations in the different tiers impact the insurance process in different ways: the tier-1 correlations will then influence an entity to seek insurance, whereas the tier-2 correlations influence the price of the premium set by the insurance company. In [46], Ogut et al. show that interdependent risks reduce the incentives of firms to invest in security and to buy insurance coverage. Our simple model (without premium discrimination) will allow to recover this result (see Section 4.3). We will show how premium discrimination can overcome this difficulty.

Kunreuther and Heal [36] consider the situation of agents faced with interdependent risks and proposes a parametric game-theoretic model for such a situation. In the model, agents decide whether or not to invest in security and agents face a risk of damage which depends on the state of other agents. They show the existence of two Nash equilibria (all agents invest or none invests), and suggest that taxation or insurance would be ways to provide incentives for agents to invest (and therefore reach the "good" Nash equilibrium), but they do not analyze the interplay between insurance and security investments. The model in [36] is extended by Hofmann in [27] to include compulsory insurance offered by a monopolistic insurer. The results show that a compulsory monopoly may lead to a higher social level of security investment if the insurer engages in premium discrimination, and that the level of investment is higher in a compulsory insurance monopoly market than in competitive insurance markets. Our work also builds on the model of [36], and considers a single insurance market. However, our work differs from [36] and [27] because it models all three desirable characteristics of an Internet-like network, namely correlated risks, interdependent agents, and a general model of a network with a flexible and controllable topology, and it derives general results about the state of the network and the behavior of the agents, with and without insurance being available.

Next, we describe the classical expected utility model for a single agent and present the standard results about premium computation and the interplay between self-protection and insurance.

3 Insurance and self-protection: basic concepts

3.1 Classical model for insurance

The classical expected utility model is named thus because it considers agents that attempt to maximize some kind of expected utility function $u[\cdot]$. In this paper, we assume that agents are rational and that they are risk averse, i.e. their utility function is concave (see Proposition 2.1 in [21]). Risk averse agents dislike mean-preserving spreads in the distribution of their final wealth. For example, consider an agent given the choice between i) a bet of either receiving \$100 or nothing, both with a probability of 50%, or ii) receiving some amount with certainty. A risk averse agent would rather accept a payoff of less than \$50 with probability 1 than the bet.

We denote by w_0 the initial wealth of the agent. The *risk premium* π is the maximum amount of money that one is ready to pay to escape a pure risk X , where a pure risk X is a centered random variable: $\mathbb{E}[X] = 0$. The risk premium corresponds to an amount of money paid (thereby decreasing the wealth of the agent from w_0 to $w_0 - \pi$) which covers the risk; hence, π is given by the following equation:

$$u[w_0 - \pi] = \mathbb{E}[u[w_0 + X]]$$

The risk premium plays a fundamental role in the economics of risk and we refer to [21] for a detailed account. We will focus in the rest of this section on the interplay between insurance and self-protection investments. To simplify our analysis, we consider simple one-period probabilistic models for the risk, in which all decisions and outcomes occur in a simultaneous instant; we do not consider dynamic aspects such as first mover advantage or the time value of money.

Each agent faces a potential loss ℓ , which we take in this paper to be a fixed (non-random) value. We denote by p the probability of loss or damage. There are two possible final states for the agent: a good state, in which the final wealth of the agent is equal to its initial wealth w_0 , and a bad state in which the

final wealth is $w_0 - \ell$. If the probability of loss is $p > 0$, the risk is clearly not a pure risk. The amount of money m the agent is ready to invest to escape the risk is given by the equation:

$$pu[w_0 - \ell] + (1 - p)u[w_0] = u[w_0 - m] \quad (2)$$

As shown by Mossin [44], we clearly have $m > p\ell$, as described on Figure 1²:




Figure 1: Computation of the risk premium: $\pi[p] = m - p\ell$.

We can actually relate m to the risk premium defined above. Note that the left hand-side of Equation (2) can be written as $\mathbb{E}[u[w_0 - p\ell - X]]$ with X defined by $\mathbb{P}(X = \ell(1 - p)) = p$ and $\mathbb{P}(X = -p\ell) = 1 - p$. Hence we have $\mathbb{E}[u[w_0 - p\ell - X]] = u[w_0 - p\ell - \pi[p]]$ where $\pi[p]$ denotes the risk premium when the loss probability equals p . Therefore:

$$m = p\ell + \pi[p].$$

The term $p\ell$ corresponds to what is referred to as the fair premium, i.e. the premium which exactly matches expected loss (for which process U defined in Equation (1) has exactly zero drift). On the left hand side of the equation, m corresponds to the maximum acceptable premium for full coverage: if an insurer makes a proposition for full coverage at a cost of φ , then the agent will accept the contract if $\varphi \leq m$. From the insurer's perspective, the premium φ depends on the distribution of the loss (here p and ℓ) and should be greater than $p\ell$ in order for the random process U defined in Equation (1) to have a positive drift. Hence the existence of a market for insuring this risk is a function of u, ℓ and p .

3.2 A model for self-protection

Investments in security involve either self-protection (to reduce the probability of a loss) and/or self-insurance (to reduce the size of a loss). For example, intrusion detection and prevention systems are mechanisms of self-protection. Denial-of-service mitigation systems, traffic engineering solutions, overprovisioning, and public relations companies are mechanisms of self-insurance (overprovisioning to reduce the impact of overloads or attacks, PR firms to reduce the impact of security attack on a company stock price with crafty messages to investors). It is somewhat artificial to distinguish mechanisms that reduce the probability of a loss from mechanisms that reduce the size of the loss, since many mechanisms do both. Nevertheless, we focus on self-protection mechanisms only and consider a very simple model for self-protection. We refer to the work of Gordon and Loeb [23] for a more elaborate model.

We first look at the problem of optimal self-protection without insurance. We denote by c the cost of self-protection and by $p[c]$ the corresponding probability of loss. We expect larger investments in self-protection to translate into a lower likelihood of loss, and therefore we reasonably assume that p is a non-increasing function of c . The optimal amount of self-protection is given by the value c^* which maximizes

$$p[c]u[w_0 - \ell - c] + (1 - p[c])u[w_0 - c]. \quad (3)$$

²The concavity of u , i.e. risk-aversion is essential here.

Note that if ℓ increases, then c^* has to increase too because the gain caused by self-protection is increased. Consider the simple case where the loss probability is either one of two values, namely $p[c] = p^+$ if $c < c_t$ or $p[c] = p^-$ if $c > c_t$, with $p^+ > p^-$. The optimization problem (3) above becomes easy to solve: indeed, the optimal expenditure is either 0 or c_t .

In the rest of the paper, we assume that the choice of an agent regarding self-protection is a binary choice: either the agent does not invest, or it invests c_t which will be denoted c for simplicity. In our case, if the agent does not invest, the expected utility is $p^+u[w_0 - \ell] + (1 - p^+)u[w_0]$; if the agent invests, the expected utility is $p^-u[w_0 - \ell - c] + (1 - p^-)u[w_0 - c]$. Using the derivation in the subsection above, we see that these quantities are equal to $u[w_0 - p^+\ell - \pi[p^+]]$ and $u[w_0 - c - p^-\ell - \pi[p^-]]$, respectively. Therefore, the optimal strategy is for the agent to invest in self-protection only if the cost for self-protection is less than the threshold

$$c < (p^+ - p^-)\ell + \pi[p^+] - \pi[p^-] =: c_1^{sp}. \quad (4)$$

Recall that $p\ell + \pi[p]$ corresponds to the amount of money the agent is willing to pay to escape a loss of probability p . Hence we can interpret Equation (4) as follows:

$$c_1^{sp} + p^- + \pi[p^-]\ell = p^+ + \pi[p^+]\ell.$$

The left hand term corresponds to the scenario where the agent invests c_1^{sp} in self-protection (and hence lower the probability of loss to p^-) and then pays $p^- + \pi[p^-]\ell$ to escape the risk. The right hand term is exactly the amount he would pay to escape the original risk of a loss of probability p^+ . Clearly the first scenario is preferred when $c < c_1^{sp}$ which corresponds exactly to Equation (4).

3.3 Interplay between insurance and self -protection

We now analyze the impact that the availability of insurance has on the level of investment in self-protection chosen by the agent.

Consider first the case when Equation (4) is satisfied, namely it is best for the agent to invest in self-protection. We assume that the agent can choose between insurance with full coverage and self-protection. Clearly if the agent chooses full coverage, he will not spend money on self-protection since losses are covered and the utility becomes $u^{fc} = u[w_0 - \varphi]$. In the case of optimal self-protection, the utility has been computed above: $u^{sp} = u[w_0 - c - p^-\ell - \pi[p^-]]$ since Equation (4) holds. Hence the optimal strategy for the agent is to use insurance if

$$c_4^{sp} := \varphi - p^-\ell - \pi[p^-] < c \quad (5)$$

Note that because of Equation (4), we must have

$$\varphi < p^+\ell + \pi[p^+]. \quad (6)$$

If Equation (4) does not hold, then it is best for the agent to not invest in self-protection, and the choice is between insurance and no self-protection. It is easy to see that if Equation (6) holds, then the premium is low enough and the optimal strategy is to pay for insurance.

The combination of insurance and self-protection raises the problem of what is referred to as moral hazard. Moral hazard occurs when agents or companies covered by insurance take fewer measures to prevent losses from happening, or maybe even cause the loss (and reap the insurance benefits from it). Indeed, if the premium does not depend on whether or not the agent invests in self-protection, then insurance becomes a negative incentive to self-protection. A known solution to the problem is to tie the premium to the amount of self-protection (and, in practice, for the insurer to audit self-protection practices and the level of care that the agent takes to prevent the loss) [17]. Note that this condition is necessary to avoid moral hazard: if the premium is not designed as above, then self-protection will be discouraged by insurance and we would observe either a large demand for insurance and a small demand for self-protection, or the converse.

A natural candidate for such a desirable premium proposed by Ehrlich and Becker [17] is the fair premium:

$$\varphi[S] = p^-\ell, \text{ and, } \varphi[N] = p^+\ell.$$

Table 1: Utility with insurance and self-protection - single user case

(I, S)	$u[w_0 - c - p^- \ell + \gamma]$
(I, N)	$u[w_0 - p^+ \ell - \gamma]$
(NI, S)	$u[w_0 - c - p^- \ell - \pi[p^-]]$
(NI, N)	$u[w_0 - p^+ \ell - \pi[p^+]]$

To agents who invest in self-protection, the insurer offers the premium $\wp[S]$ and to agents who do not invest in self-protection, he offers the premium $\wp[N]$. Since $p^- \leq p^+$, with such a choice, the price of insurance is negatively related to the amount of self-protection. With this premium, it is proved in [17] that insurance can co-exist with an incentive to invest in self-protection in some cases (if the probability of loss is not very small).

We will show that, even if the fair premium is negatively related to the amount spent in self-protection, it is not always sufficient for insurance to be an incentive for self-protection when risks are interdependent. In order to raise the social level of self-protection, the insurer may engage in premium discrimination. In particular, he may design different contracts for different risk types, relying on the policyholders' categorization: he may offer a premium rebate for low risk agents, and/or he may impose a premium loading for high risk agents and let agents voluntarily decide whether or not to invest in self-protection. The sequence of the considered game between the insurer and its customers may then be seen as follows: at a first stage, the insurer offers appropriate contracts including a premium loading and/or rebate on fair premiums. At a second stage, the customers choose a contract and decide simultaneously whether or not to invest in prevention. To agents who do not invest in prevention, the insurer may offer a premium $\wp[N] + \gamma$, where $\gamma \geq 0$ denotes a premium penalty (loading). To agents who invest in prevention, the insurer may offer a premium $\wp[S] - \gamma$, where γ denotes a premium rebate.

The utility for all possible cases is summarized in Table 1. The first column denotes the choice made by an agent. It is denoted by the pair (U, V) , where $U = I$ means that the agent pays for insurance and $U = NI$ otherwise, and $V = S$ means that the agent invests in self-protection and $V = N$ otherwise.

Note that for any non-negative value of γ , the strategy (I, S) always dominates the strategy (NI, S) . Now for (I, S) to dominate (I, N) , we need

$$c < (p^+ - p^-)\ell + 2\gamma.$$

For (I, S) to dominate (NI, N) , we need

$$c < (p^+ - p^-)\ell + \gamma + \pi[p^+].$$

The results are summarized in Figure 2.




Figure 2: Full coverage vs self-protection - single user case

The grayed area corresponds to the space where the parameter c is such that not investing in self-protection is optimal (N). Each row corresponds to a different case:

- The first row NI corresponds to the case when no insurance is available;
- The second row I corresponds to the case when a full coverage insurance is available with premium \wp satisfying Equation (6);

- The third row $I, \gamma = 0$ corresponds to the case when a full coverage insurance is available with premium defined as above with $\gamma = 0$ (as in [17]);
- The fourth row $I, \gamma > 0$ is the same as the row above but with a strictly positive value of γ .

The pair (I, S) in row 3, for example (resp. the pair (NI, S) in row 2) means that insurance and self-protection (resp. no insurance and self-protection) is the optimal strategy for those values of c . We have

$$\begin{aligned}
c_1^{sp} &= (p^+ - p^-)\ell + \pi[p^+] - \pi[p^-], \\
c^{sp}[\gamma] &= (p^+ - p^-)\ell + \gamma + \min(\gamma, \pi[p^+]), \\
c_3^{sp} &= (p^+ - p^-)\ell, \\
c_4^{sp} &= \wp - p^- \ell - \pi[p^-].
\end{aligned}$$

Note in particular that as soon as $\gamma > (\pi[p^+] - \pi[p^-])/2$, then we have $c^{sp}[\gamma] > c_1^{sp}$, in which case, insurance is an incentive for self-protection. This concludes the description of results from classical insurance theory. Next, we consider a 2-agent model (the first step towards the general network model), with correlated and interdependent risks. We first describe known results in the absence of insurance, then present our new results, with insurance available to agents.

4 Interdependent security and insurance: the 2-agent case

Recall that interdependent risks are risks that depend on the behavior of other entities in the network (e.g. whether or not they invested in security solutions to handle their risk). In the presence of interdependent risks, the reward for a user investing in self-protection depends on the general level of security in the network.

4.1 Interdependent risks for 2 agents

Reference [36] was the first to introduce a model for interdependent security (IDS), specifically a model for two agents faced with interdependent risks, and it proposed a parametric game-theoretic model for such a situation. In the model, agents decide whether or not to invest in security and agents face a risk of damage which depends on the state of other agents. As in Section 3 above, the decision is a discrete choice: an agent either invests or does not invest in self-protection. We assume that loss can happen in two ways: it can either be caused directly by an agent (direct loss), or indirectly via the actions of other agents (indirect loss). We assume that the cost of investing in self-protection is c , and that a direct loss can be avoided with certainty when the agent has invest in self-protection.

The cost of protection should not exceed the expected loss, hence $0 \leq c \leq p\ell$. Four possible states of final wealth of an agent result: without protection, the final wealth is w_0 in case of no loss and $w_0 - \ell$ in case of loss. If an agent invests in protection, its final wealth is $w_0 - c$ in case of no loss and $w_0 - c - \ell$ in case of loss.

Consider now a network of 2 agents sharing one link. There are four possible states denoted by (i, j) , where $i, j \in \{S, N\}$, i describes the decision of agent 1 and j the decision of agent 2, S means that the agent invests in self-protection, and N means that the agent does not invest in self-protection. We examine the symmetric case when the probability of a direct loss is p for both agents, where $0 < p < 1$. Knowing that one agent has a direct loss, the probability that a loss is caused indirectly by this agent to the other is q , where $0 \leq q \leq 1$. Hence q can be seen as a probability of contagion. To completely specify the model, we assume that direct losses and contagions are independent events. The matrix $p(i, j)$ describing the probability of loss for agent 1, in state (i, j) , is given in Table 2.

The simplest situation of interdependent risks, involving only two agents, can be analyzed using a game-theoretic framework. We now derive the payoff matrix of expected utilities for agents 1 and 2. If both agents invest in self-protection, the expected utility of each agent is $u[w_0 - c]$. If agent 1 invests in self-protection (S) but not agent 2 (N), then agent 1 is only exposed to the indirect risk pq from agent 2. Thus the expected utility for agent 1 is $(1 - pq)u[w_0 - c] + pq u[w_0 - c - \ell]$ and the expected utility for

Table 2: Probability of states

	S	N
S	$p[S, S] = 0$	$p[S, N] = pq$
N	$p[N, S] = p$	$p[N, N] = p + (1 - p)pq$

Table 3: Expected payoff matrix for agent 1

	agent 2: S	agent 2: N
agent1 : S	$u[w_0 - c]$	$(1 - pq)u[w_0 - c] + pqu[w_0 - c - \ell]$
agent1 : N	$(1 - p)u[w_0] + pu[w_0 - \ell]$	$pu[w_0 - \ell] + (1 - p)(pqu[w_0 - \ell] + (1 - pq)u[w_0])$

agent 2 is $(1 - p)u[w_0] + pu[w_0 - \ell]$. If neither agent invests in self-protection, then both are exposed to the additional risk of contamination from the other. Therefore, the expected utilities for both agents are $pu[w_0 - \ell] + (1 - p)(pqu[w_0 - \ell] + (1 - pq)u[w_0])$. Table 3 summarizes these results and gives the expected utility of agent 1 for the different choices of the agents.

Assuming that both agents decide simultaneously whether or not to invest in self-protection, there is no possibility to cooperate. For investment in self-protection (S) to be a dominant strategy, we need

$$\begin{aligned} u[w_0 - c] &\geq (1 - p)u[w_0] + pu[w_0 - \ell] \text{ and} \\ (1 - pq)u[w_0 - c] + pqu[w_0 - c - \ell] &\geq \\ pu[w_0 - \ell] + (1 - p)(pqu[w_0 - \ell] + (1 - pq)u[w_0]) \end{aligned}$$

With the notations introduced earlier, the inequalities above become:

$$\begin{aligned} c &\leq p\ell + \pi[p] =: c_1, \\ c &\leq p(1 - pq)\ell + \pi[p + (1 - p)pq] - \pi[pq] =: c_2. \end{aligned}$$

In most practical cases, one expects that $c_2 < c_1$, and the tighter second inequality reflects the possibility of damage caused by other agent. Therefore, the Nash equilibrium for the game is in the state (S, S) if $c \leq c_2$ and (N, N) if $c > c_1$. If $c_2 < c \leq c_1$, then both equilibria are possible and the solution to the game is indeterminate. More precisely, the situation corresponds to a coordination game. Overall, we have the following:

- if $c < c_2$: the optimal strategy is to invest in self-protection;
- if $c_2 < c < c_1$: if the other user in the network do invest in self-protection, then the optimal strategy is to invest in self-protection;
- if $c_1 < c$: then the optimal strategy is to not invest in self-protection.

4.2 IDS and mandatory insurance

We now build on the model and the results above and introduce our more general model in which insurance is available to the agents (the ability to self-protect remaining available, of course). We assume that a full coverage insurance is mandatory. As noted in Section 3.3, if we want to avoid a moral hazard problem, the insurance premium has to be tied to the amount spent on self-protection. Note that the probability of loss for agent 1 depends on the choice made by agent 2, however it seems necessary (at least from a practical point of view) to link the premium applied to agent 1 to the behavior of agent 1 only. A possible choice (which is profit-making for the insurance) is to choose for each decision of the agent the fair 'worst case' premium as follows,

$$\wp[S] = pq\ell, \quad \wp[N] = (p + (1 - p)pq)\ell.$$

In this case the payoff for the agent is deterministic: if it chooses S , the payoff is $u[w_0 - c - pq\ell]$; if it chooses N , the payoff is $u[w_0 - (p + (1 - p)pq)\ell]$. Hence the dominant strategy is to invest in self-protection only if

$$c < p(1 - pq)\ell =: c_3 < c_2.$$

Table 4: Expected payoff matrix with insurance and self-protection

	agent 2: S	agent 2: N
(I, S)	$u[w_0 - c - pq\ell + \gamma]$	
(I, N)	$u[w_0 - (p + pq(1 - p))\ell - \gamma]$	
(NI, S)	$u[w_0 - c]$	$(1 - pq)u[w_0 - c] + pqu[w_0 - c - \ell]$
(NI, N)	$(1 - p)u[w_0] + pu[w_0 - \ell]$	$pu[w_0 - \ell] + (1 - p)(pqu[w_0 - \ell] + (1 - pq)u[w_0])$

As in the single-agent case, we see that even if the premium is related to the amount spent on self-protection, insurance is a negative incentive for protection. To correct this effect, we apply the same strategy as in the single-agent case, namely we engage in premium discrimination. Let γ denote the premium rebate for agents investing in security and the premium penalty for agents not investing. Clearly, in our situation, the new condition for S to be the dominant strategy becomes:

$$c < p(1 - pq)\ell + 2\gamma =: c_3[\gamma].$$

In particular for $2\gamma = p^2q\ell$, we have $c_3[\gamma] = c_1$ and then for any $c < c_1$, the strategy S is dominant (whereas coordination was required in absence of insurance). Note that we have assumed a symmetric penalty and rebate but our result easily extends to the general case.

4.3 IDS and full coverage insurance

We now consider the situation where the choice is left to the agent as to whether to invest in self-protection and/or in a full coverage insurance. We assume that the premiums are those given above (with penalty/rebate). We summarize the payoff for agent 1 in Table 4, depending on the investment of agent 2 and for the four possible choices of the agent (notations are the same as in Section 3.3). We denote

$$c_4[\gamma] := p(1 - pq)\ell + \pi[p + (1 - p)pq] + \gamma.$$

Let us examine the situation depending on the behavior of agent 2. If agent 2 invests in self-protection (denoted by S_2), then for $c < c_1$, agent 1 chooses to invest in self-protection also and not otherwise.

Consider now the case when agent 2 does not invest in self-protection (denoted by N_2). Then if $c < \min(c_3[\gamma], c_4[\gamma]) := c[\gamma]$, the optimal strategy is (I, S) . Note that we have $c_4[\gamma] \geq c_2$ for all values of γ and we proved above that we can choose γ such that $c_3[\gamma] \geq c_2$. Therefore it is possible to tune γ such that $c[\gamma] \geq c_2$.




Figure 3: Full coverage vs self-protection

Figure 3 summarizes the results. As in Figure 2, the grayed area corresponds to the space where the parameter c is such that not investing in self-protection is optimal (N). Each row corresponds to a different case:

- The first row NI corresponds to the case when no insurance is available;
- The second row I corresponds to the case when a full coverage insurance is compulsory;
- The third row I, γ, S_2 corresponds to the case when a full coverage insurance is available with premium defined as above with $\gamma > 0$ and agent 2 is in state S ;

- The fourth row I, γ, N_2 is the same as the row above but when agent 2 is in state N .

The pair (I, S) in row 4, for example (resp. the pair (NI, S) in row 3) means that the combination of insurance and self-protection (resp. no insurance and self-protection) is the optimal strategy for those values of c . We have (in decreasing order)

$$\begin{aligned} c_1 &= p\ell + \pi[p], \\ c[\gamma] &= p(1-pq)\ell + \gamma + \min(\pi[p + (1-p)pq], \gamma), \\ c_2 &= p(1-pq)\ell + \pi[p + (1-p)pq] - \pi[pq], \\ c_3 &= p(1-pq)\ell. \end{aligned}$$

Note in particular that when insurance with discrimination is available, (S, S) becomes a Nash equilibrium for $c < c[\gamma]$ with $c[\gamma] > c_2$ for well-chosen values of γ . In such a case, insurance is an incentive to self-protection. However, if insurance is available at a fair premium, without discrimination, (i.e. $\gamma = 0$), then we see that $c[0] = c_3 < c_2$ and insurance is not anymore an incentive to self-protection.

The main features present in the single-agent (Figure 2) are also present in the 2-agent case (Figure 3). However a new feature comes into play because of the interdependent risks, namely the existence of a new threshold c_2 which takes into account the externality modeled by the possible contagion via the other agent. We see that the externalities due to the interdependent risks tend to lower the incentive for investing in self-protection (as shown in [46]). However, we also see that the effect of the insurance (with discrimination) is unaffected by these interdependent risks. As a result the relative efficiency of insurance is higher in the presence of externalities.

Next, we extend the results of this section to the general case of a network of n users.

5 Interdependent security and insurance on a network

Many phenomena in the Internet can be modeled using epidemic spreads through a network, e.g. the propagation of worms, of email viruses, of alerts and patches, of routing updates, etc. (e.g. see [56, 54] for models of worm propagation). As a result, there is now a vast body of literature on epidemic spreads over a network topology from an initial set of infected nodes to susceptible nodes (see for example [20]). The 2-agent model introduced in the previous section, although very basic, fits in that framework: the probability for an agent to be infected initially is p and the probability of contagion is q . It is then natural to consider the following extension: agents are represented by vertices of a graph $G = (V, E)$, and

- if an edge $(i, j) \in E$ then contagion is possible between agents i and j with probability q ; otherwise the probability of contagion is zero;
- if agent i invests in protection, no direct loss can occur; otherwise direct loss occurs with probability p ;
- the contagion process of agent i is independent of the process of agent j and independent of the direct loss process (characterized by p).

As in the previous section, we are considering one-period model. The quantity of interest here is the value of the damages due to the epidemics. We assume that the damage caused by the epidemics is ℓ for all agents that have been infected.

The topology of the underlying graph G is arbitrary. Note that G might not correspond to a physical network. For example, when modeling the spread of email viruses, we might choose a graph which reflects the social network of the email users. When modeling insurance against BGP router failures, we might choose a complete graph; indeed, BGP routers belonging to the top level ASes of the Internet form a completely connected graph, and internal BGP routers are often organized in a set of completely connected route reflectors - thus, the behavior of routers failing and recovering is, in a first approximation, modeled as the spread of an epidemic on a complete graph [14].

In the rest of this section, we consider two important classes of topologies: the complete graph and the star-shaped graph. The study of star-shaped networks is of interest for several reasons. First, star-shaped

networks exhibit a new tipping point phenomenon not observed in fully connected networks. Also, the spreading behavior of a large class of power law graphs, of particular interest given their relevance to Internet topology graphs [16], is determined by the spreading behavior of stars embedded within them [20].

5.1 The complete graph network

We assume here that G is a complete graph with n vertices, namely a graph with an edge between each pair of nodes. By symmetry, it is possible to define P_k^S , the probability that an agent investing in security experiences a loss when k users (among the $n - 1$ remaining ones) also invest in security. Similarly, we define P_k^N to denote the probability that a user not investing in security experiences a loss when k other users invest in security. Then we define:

$$c_k^n = P_k^N \ell + \pi[P_k^N] - P_k^S \ell - \pi[P_k^S]. \quad (7)$$

We have of course $P_k^N \geq P_k^S$ and we assume that the utility function u (which defines the function π) is such that $c_{k+1}^n \geq c_k^n$ for all $0 \leq k \leq n - 1$. Note that in the single user case, $n = 1$, we have $c_0^1 = c_1^{sp}$ defined in Equation (4). In the 2-user case, we have $c_1^2 = c_1$ and $c_0^2 = c_2$ defined in Section 4.1.

Results of Section 4 extend in a straightforward manner to the n -users case as follows:

- if $c < c_0^n$: the optimal strategy is to invest in self-protection;
- if $c_{k-1}^n < c < c_k^n$: if at least k users in the network do invest in self-protection, the optimal strategy is to invest in self-protection;
- if $c_{n-1}^n < c$: the optimal strategy is to not invest in self-protection.

It is natural to define the following function:

$$k^n[c] = \inf\{k, c_k^n > c\}.$$

$k^n[c]$ is an important threshold value, because of the following:

- if the number of initial users investing in self-protection is less than $k^n[c]$, then all users will chose not to invest in self-protection;
- if the number of initial users investing in self-protection is greater than $k^n[c]$, then all users will chose to invest in self-protection.

Concerning the effect of an insurance, we only consider the case where the insurance company engages in premium discrimination. It is then easy to extend the results above with the function $c^n[\gamma]$ such that if $c < c^n[\gamma]$, then the optimal strategy is to invest in self-protection regardless of the behavior of the other users. Furthermore, $c^n[\gamma]$ is a non-decreasing function of γ that tends to infinity as γ tends to infinity.

In summary, we have the following simple situation: in presence of insurance, the optimal strategy for all users is to invest in self-protection as soon as the cost of self-protection is low enough $c < c^n[\gamma]$.

The situation is simple, but artificially so, because we are considering a purely symmetric case. Let us now consider the more general case of *heterogeneous users*, when the cost of self-protection is different for different users (but the effect of self-protection is not changed). Intuitively, users with low cost will tend to invest in prevention while those with high cost will not. We now derive the threshold \hat{c} for which users with cost less than \hat{c} invest in self-protection whereas others do not. We denote by $F^n[c]$ the fraction of users with self-protection cost lower than c . Let s_j denote the different possible values for the cost of self-protection. The function F^n is piecewise constant and increases at each s_j by the fraction of nodes having a cost of s_j .

Consider now the following dynamic process where all the users of the network are initially in state (N) , i.e. they have not invested in self-protection. First consider the users with minimal cost, say s_0 . If $s_0 < c_0^n$, then $nF^n[s_0]$ users switch and invest in self-protection. If $s_0 > c_0^n$, all users stay in state (N) and the process terminates. Next, consider the users still in state (N) with minimal cost s_1 . If

$s_1 < c_{nF_n[s_0]}^n$, then all those users will switch and invest in self-protection. Note that the condition above can be written as $k^n[s_1] < nF_n[s_0]$. Iterating the procedure, we see that the threshold is characterized by the following equation

$$\hat{c} = \min\{s_{j-1}, F^n[s_{j-1}] < \frac{k^n[s_j]}{n}\}. \quad (8)$$

In order to analyze the impact of insurance on the dynamics of the process above, we approximate the n users by a continuum of heterogeneous users. Showing that this mean-field approximation is appropriate for large values of n is outside the scope of this paper and requires a scaling of the probabilities p and q as n tends to infinity. However we present the following heuristic argument. We denote by $F[c]$ the distribution function of the users and by $k[c]$ the limit of $k^n[c]/n$, both of which now continuous. Then Equation (8) reduces to

$$\hat{c} = \min\{c, F[c] = k[c]\}.$$

When adding assurance, the same argument as above holds, but this time we can start with an initial condition where all users with cost less than $c[\gamma]$ invest in self-protection. Hence the final equilibrium will be given by

$$\hat{c}[\gamma] = \min\{c > c[\gamma], F[c] = k[c]\}.$$

Note that for any value of $\gamma \geq 0$, we have $\hat{c}[\gamma] \geq \hat{c}$ which shows that *more users choose to invest in self-protection in presence of insurance*. Furthermore, if

$$F[c] = k[c] \quad (9)$$

has only one solution, then $\hat{c}[\gamma] = \max\{c[\gamma], \hat{c}\}$.

The results above show that insurance increases the adoptability of self-protection investments for all users in the network. We finish this section by showing that the increase in adoptability can be quite dramatic, non-linearly so as a function of γ .

Assume now that the population is divided into classes of users with roughly the same cost for self-protection and consider the case when users corresponding to the class with the smallest cost invest in self-protection. If the size of that population (of users in the class with the smallest cost) is small, it might not be sufficient to stimulate the second class³ to invest in self-protection too. Then the dynamics of the 'contagion process' for self-protection described earlier is stopped and only a small fraction of the total population has invested in security in the end. It turns out that insurance can be of significantly help to boost the contagion process, as we explain next. Note that the function F is approximately a step function and Equation (9) might have more than one solution, see Figure 4. The scenario described above corresponds to the case when the system is stuck at the low value \hat{c} . We see that if we tune the parameter γ in order for $c[\gamma]$ to reach the second fixed point, then the system will naturally increase its level of self-protection up to the next fixed point $\hat{c}[\gamma] \gg c[\gamma] > \hat{c}$ as described on Figure 4. In other words, insurance gives exactly the right incentive to a small portion of the population that would have not invested without insurance, so that the switch to self-protection of that fraction of the population induces a larger fraction of the population to invest also. **In summary, insurance provides incentives for a small fraction of the population to invest in self-protection, which in turn induces the rest of the population to invest in self-protection as well, leading to the desirable state where all users in the network are self-protected. Furthermore, the parameter γ provides a way to multiply the benefits of insurance, by lowering the initial fraction of the self-protected population needed to reach the desirable state.**

5.2 The star-shaped network

Consider a star-shaped network, with $n + 1$ nodes, where the only edges are $(0, i)$, with $i = 1, \dots, n$. The same analysis as in previous section applies but we have to deal separately with the root and the leaves.

³i.e. the next set of users with the second smallest cost of self-protection

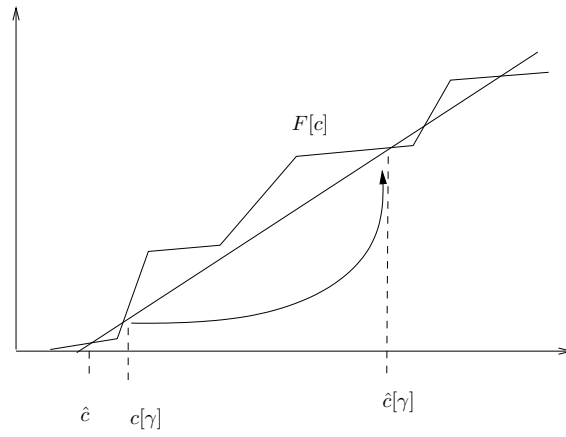


Figure 4: Non linearity of insurance

First consider the root. The probability of a loss when exactly k leaves invest in self-protection is given (depending on the state (S) or (N) of the root) by

$$\begin{aligned} P_k^N &= p + (1-p)(1 - (1-pq)^{n-k}), \\ P_k^S &= (1 - (1-pq)^{n-k}). \end{aligned}$$

One can then do the same analysis as in previous section and compute the function $k^n[c]$ that would give the threshold for the number of leaves required to invest in self-protection in order for the root to also invest. Note that as n tends to infinity, the probability of loss tends to one as soon as the number of leaves not investing in self-protection tends to infinity. In this case, the agent at the root is sure to be contaminated by a leaf regardless of its choice regarding investment in self-protection. As a result, it will not invest in self-protection.

We next consider the leaves. An important remark is that for a leaf to be infected the root must also be infected. First assume that the root is in state (N). The probabilities of loss when there are k other leaves investing in self-protection are:

$$\begin{aligned} \tilde{P}_k^N &= p + (1-p)qP_k^N[n], \\ \tilde{P}_k^S &= qP_k^N[n], \end{aligned}$$

where $P_k^N[n]$ is the quantity computed above but for a network of size n . Now assume that the root is in state (S), the corresponding probabilities are given by (with the same notations)

$$\begin{aligned} \bar{P}_k^N &= p + (1-p)qP_{k-1}^S[n], \\ \bar{P}_k^S &= qP_{k-1}^S[n]. \end{aligned}$$

It is easy to see that $\bar{P}_k^N - \bar{P}_k^S \leq \tilde{P}_k^N - \tilde{P}_k^S$, and as a result, we have

$$\bar{c}_k^{n+1} \leq \tilde{c}_k^{n+1}, \quad (10)$$

where the parameters \bar{c}_k^{n+1} and \tilde{c}_k^{n+1} are defined as in Equation (7) with the appropriate probabilities. The incentive for a leaf to invest in self-protection is higher when the root already invested in self-protection. We observe a tipping point phenomenon. More generally, we expect that nodes with low connectivity (i.e. low degree) will imitate the node with the highest connectivity they are connected to. The heuristic argument (which is captured in our model by (10)) is that the node with the largest connectivity an agent is connected to will be the main source of contagion (in term of probability). Hence, if that node invests in self-protection, it substantially decreases the probability of contagion of the agent and, as in the 2-agent case, that action increases the reward of investing in self-protection. In such a context, insurance could act as an incentive for highly connected agents to invest in self-protection and then trigger a cascade of adoption of self-protection. A precise analysis of this phenomena is left for future research.

6 Discussion

The results presented in this paper show that insurance provides significant benefits to a network of users facing correlated, interdependent risks. Essentially, insurance is a powerful mechanism to promote network-wide changes and lead all users of the network to the desirable state where they all invest in self-protection. The benefits of insurance are such that we believe that the development of insurance products and markets, and the large scale deployment of insurance in the Internet is likely, if not inevitable.

However, we have found that mentioning "Internet insurance" rapidly attracts comments about the uniqueness of the Internet environment, and in particular questions around the **estimation of damages**. The assumption is that estimating damages in the Internet is so difficult and fraught with peril that insurance is not inevitable at all, but rather destined to remain a niche or an oddity. We first note that reliably estimating damages is indeed an important task because it controls the profit (or the ruin) of the insurer and the incentives for agents to invest in self-protection. Also, it is true that quantifying risks for a good or an optimal premium value is difficult because the assets to be protected are intangible (such as a company stock price), because damages might be visible only long after a threat or an attack was identified (e.g. "easter egg" with timed virus or exploit in a downloaded piece of software), because risk changes can occur quickly (zero day attacks), and because evaluating the insurability (and the level of protection) of new and existing customers is likely to be a complex and time intensive task. However, the insurance industry has been dealing with those problems for decades or centuries in other areas of life - if warships can be insured in time of war (as indeed they can), it is difficult to argue convincingly that Internet risks and damages absolutely cannot be insurable.

Questions about damage estimation might also be the wrong questions. A better question might be **how to help insurers do a better job**, i.e. how the current Internet might be used to help insurers do a better job of estimating damages, and how to evolve the Internet or create a new design that will make that job even easier. One way suggested by the discussion above on estimating damages would be to develop metrics and techniques for that purpose. Another, related way is to develop metrics for the security related issues of interest. Some interesting propositions have been made in that sense, for example the "cost to break" metric described in [49], but we believe this is an important area ripe for further research (see also [3]). Note that metrics of interest are not limited to "core security" metrics such as cost to break, but need to be developed for all relevant activities facing threats and risks; for example, metrics quantifying risks and damages to insure against BGP router failures (mentioned in Section 5).

The deployment of insurance raises **architectural issues**. In particular, insurance relies heavily on authenticated, audited, or certified assessments of various kinds to avoid fraud and other issues such as the moral hazard examined earlier in the paper. This argues, along with security metrics, for effective and efficient ways to measure and report those metrics. It might also require better traceability of events. But it will certainly impact other mechanisms and protocols in other, subtle ways. Consider for example a peering point between operators, some of which are insured, others of which are not. It is very reasonable to imagine that, in such a situation, policies would be developed to route traffic from insured peers (or neighbors in general) differently than traffic from un-insured peers - a latter-day QoS routing (where QoS means Quality of Security, of course).

Overall, we believe that Internet insurance, in addition to providing the benefits shown in the paper, offers a fertile area of reflection and research. It is a timely area, as well, given the recent activities around clean-slate Internet design. **We propose to add to the slate a broader definition of risk management, which includes the transfer of risk in addition to only the mitigation of risk, and explore the benefits and consequences of that broader definition.**

7 Conclusion

One of our main contributions in this paper is to develop and solve simple models that explain why economically rational entities would prefer a relatively insecure system to a more secure one, that show that the adoption of security investments follows a threshold or tipping point dynamics, and that insurance is a powerful incentive mechanism to "push the mass of users over the threshold". Our second main contribution is to ask the question: if economics plays an essential role in the deployment of security technologies then why deny ourselves the use of economics tools? Our purpose here is not to shift the

problem of network security to the marketplace but to give a new perspective on Internet security. Finally, we argue that network algorithms and network architecture might be designed or re-evaluated according to their ability to help implement desirable economic policies (such as the deployment of insurance) and help achieve desirable economic goals.

References

- [1] American International Group, Inc. netAdvantage <http://www.aignetadvantage.com/>
- [2] R. Anderson and T. Moore. The economics of information security: A survey and open questions. *Science*, vol. 314, pp. 610-613, Oct. 2006.
- [3] J. Aspnes et al. Towards better definitions and measures of Internet security. *Proc. Workshop on Large-Scale-Network Security and Deployment Obstacles*, Landsdowne, VA, March 2003.
- [4] D. A. Barnes. Deworming the internet. *Texas Law Review*, 83(1), 2004. Available at SSRN:<http://ssrn.com/abstract=622364>.
- [5] J. Bolot and M. Lelarge. A New Perspective on Internet Security using Insurance. *INFOCOM 08 (Mini-Conference)*.
- [6] R. Böhme. Cyber-insurance revisited. *Proc. of Workshop on the Economics of Information Security (WEIS)*, 2005.
- [7] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. *Proc. of Workshop on the Economics of Information Security (WEIS)*, 2006.
- [8] L. J. Camp and C. Wolfram. Pricing security. *Proc. CERT Information Survivability Workshop*, Boston, MA, pp. 24-26, Oct. 2000.
- [9] H. Chan, D. Dash, A. Perrig, H. Zang. Modeling adoptability of secure BGP protocols. *Proc. ACM Sigcomm 06*, Pisa, Italy, Sept. 2006.
- [10] W. R. Cheswick, S. Bellovin, A. Rubin, *Firewalls and Internet Security: Repelling the Wily Hacker*, 2nd Ed., Addison-Wesley, 2003.
- [11] P. Chen et al. Software diversity for information security. *Proc. WEIS 2005*, Harvard, MA, June 2005.
- [12] D. Clark. The design philosophy of the DARPA Internet protocols. *Proc. ACM Sigcomm 88*, Stanford, CA, Aug 1988.
- [13] D. Clark, J. Wroclawski, K. Sollins, R. Braden. Tussle in cyberspace: defining tomorrow's internet. *Proc. ACM Sigcomm 02*, Pittsburgh, PA, Aug. 2002.
- [14] E.G. Coffman Jr., Z. Ge, V. Misra and D. Towsley. Network resilience: exploring cascading failures within BGP. *Proc. 40th Annual Allerton Conference on Communications, Computing and Control, October 2002*.
- [15] <http://www.counterpane.com/pr-lloydssl.html>
- [16] J. C. Doyle et al. The "robust yet fragile" nature of the Internet. *Proc. Nat. Acad. Sciences*, vol. 102, no. 41, Oct 2005.
- [17] I. Ehrlich and G. S. Becker. Market insurance, self-insurance, and self-protection. *The Journal of Political Economy*, 80(4):623-648, 1972.
- [18] Elsevier Ed. *Insurance: Mathematics and Economics*.
- [19] FIND: NSF Future INternet Design program area.
- [20] A. Ganesh, L. Massoulie, D. Towsley. The effect of network topology on the spread of epidemics. *Proc. IEEE Infocom 2005*, Miami, FL, March 2005.

- [21] C. Gollier. *The Economics of Risk and Time*. MIT Press, 2004.
- [22] J. Gong and P. Srinagesh. The economics of layered networks. *Internet Economics*, MIT Press, Cambridge, MA, 1997.
- [23] L. Gordon and M. Loeb. The economics of information security investment. *ACM Trans. Inf. Syst. Secur.* 5, 4 (Nov. 2002), 438-457.
- [24] L. Gordon and M. Loeb. *Managing Cybersecurity Resources*. McGraw-Hill, Sept. 2005.
- [25] L. Gordon, M. Loeb and T. Sohail. A framework for using insurance for cyber-risk management. *Commun. ACM*, 46(3), pp. 81-85, 2003.
- [26] H. Herath and T. Herath. Cyber-Insurance: Copula Pricing Framework and Implications for Risk Management. *Proc. of Workshop on the Economics of Information Security (WEIS)*, 2007.
- [27] A. Hofmann. Internalizing externalities of loss prevention through insurance monopoly. *Proc. Annual Meeting of American Risk and Insurance Association*, Washington DC, Aug 2006.
- [28] IEEE Symposium on Security and Privacy www.ieee-security.org/TC/SP-Index.html
- [29] Internet Measurement Conference www.imcconf.net
- [30] J. Jung, V. Paxson, A. Berger, H. Balakrishnan. Fast portscan detection using sequential hypothesis testing In *Proc. IEEE Symp. Security and Privacy*, 2004.
- [31] G. Davie, M. Hardt, F. Kelly. Network dimensioning, service costing, and pricing in a packet switched environment. In *Telecommunications Policy*, vol. 28, pp. 391-412, 2004.
- [32] M. Kearns and L. E. Ortiz. Algorithms for interdependent security games. In *Advances in Neural Information Processing Systems*, S. Thrun, L. K. Saul and B. Schoikopf, Eds., MIT Press, Cambridge, 2004.
- [33] J. Kesan, R. Majuca, and W. Yurcik. The economic case for cyberinsurance. In *Securing Privacy in the Internet Age*, A. Chander et al., Eds., Stanford University Press, 2005.
- [34] J. Kesan, R. Majuca, and W. Yurcik. Cyberinsurance as a market-based solution to the problem of cybersecurity: a case study. *Proc. WEIS 2005*, Harvard, MA, June 2005.
- [35] L. Kleinrock. Research areas in computer communications. *Computer Comm. Review*, v. 4, no. 3, July 1974.
- [36] H. Kunreuther and G. Heal. Interdependent security: the case of identical agents. *Journal of Risk and Uncertainty*, 26(2):231-249, 2003.
- [37] M. Lelarge and J. Bolot. Network Externalities and the Deployment of Security Features and Protocols in the Internet. *ACM SIGMETRICS 08*.
- [38] M. Lelarge and J. Bolot. A Local Mean Field Analysis of Security Investments in Networks. *ACM NETECON 08*.
- [39] C. Lai et al. Endorsments, licensing, and insurance for distributed systems services. *Proc. 2nd ACM Conf. Computer and Comm. Security (CCS)*, Fairfax, VA, Nov. 1994.
- [40] J. MacKie-Mason and H. Varian. Pricing the Internet. in B. Kahin and J. Keller, Eds., *Public Access to the Internet*, MIT Press, 1995.
- [41] R. P. Majuca, W. Yurcik, and J. P. Kesan. The evolution of cyberinsurance. *Information Systems Frontier*, 2005.
- [42] T. Mikosh. *Non-Life Insurance Mathematics: An Introduction with Stochastic Processes*. Springer, June 2006.
- [43] Network and Distributed Systems Symposium <http://www.isoc.org/isoc/conferences/ndss/>
- [44] J. Mossin. Aspects of rational insurance purchasing. *Journal of Political Economy*, 76:553-568, 1968.

- [45] A. Odlyzko. Economics, psychology, and sociology of security. *Proc. Financial Cryptography 2003*, R. N. Wright, Ed., LNCS #2742, Springer, Apr. 2003.
- [46] H. Ogut, N. Menon and S. Raghunathan. Cyber Insurance and IT Security Investment: Impact of Interdependent Risk. *Workshop Economics Info. Sec. (WEIS)*, 2005.
- [47] A. Ozment and S. Schechter. Bootstrapping the adoption of Internet security protocols. *Workshop Economics Info. Sec.*, Cambridge, June 2006.
- [48] S. Saniford, D. Moore, V. Paxson, N. Weaver. The top speed of flash worms. *Proc. ACM Workshop Rapid Malcode WORM'04*, Fairfax, VA, Oct 2004.
- [49] S. Schechter. Quantitatively differentiating system security. *Workshop Economics Info. Sec.*, Berkeley, May 2002.
- [50] B. Schneier. Insurance and the computer industry. *CACM*, vol. 44, no. 3, March 2001.
- [51] B. Schneier. Computer security: It's the economics, stupid. *Proc. WEIS 2002*, Berkeley, CA, May 2002.
- [52] S. Shenker, D. Clark, D. Estrin, S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM CCR*, vol. 26, pp. 19-43, Apr. 1996.
- [53] H. Varian, J. Farrell, C. Shapiro. *The Economics of Information Technology*. Cambridge University Press, Dec. 2004.
- [54] M. Vojnovic and A. Ganesh. On the race of worms, alerts and patches. *Proc. ACM Workshop on Rapid Malcode WORM05*, Fairfax, VA, Nov. 2005.
- [55] N. Weaver and V. Paxson. A worst-case worm. *Proc. 3rd Workshop Economics Info. Sec.*, Univ. Minnesota, May 2004. See web site for opinion by S. Saniford.
- [56] C. Zou, W. Gong, D. Towsley. Code Red worm propagation modeling and analysis. *Proc. 9th ACM Conf. Computer Comm. Security CCS'02.*, Washington, DC, Nov 2002.