

Security Games with Market Insurance

Benjamin Johnson¹, Rainer Böhme², and Jens Grossklags³

¹ Department of Mathematics, UC Berkeley

² Department of Information Systems, University of Münster

³ College of Information Sciences and Technology, Penn State University

`benjamin@math.berkeley.edu`,

`rainer.boehme@wi.uni-muenster.de`,

`jensg@ist.psu.edu`

Abstract. Security games are characterized by multiple players who strategically adjust their defenses against an abstract attacker, represented by realizations of nature. The defense strategies include both actions where security generates positive externalities and actions that do not. When the players are assumed to be risk averse, market insurance enters as a third strategic option. We formulate a one-shot security game with market insurance, characterize its pure equilibria, and describe how the equilibria compare to established results. Simplifying assumptions include homogeneous players, fair insurance premiums, and complete information except for realizations of nature. The results add more realism to the interpretation of analytical models of security games and might inform policy makers on adjusting incentives to improve network security and foster the development of a market for cyber-insurance.

Keywords: Game theory, Security, Externalities, Protection, Self-insurance, Market insurance.

1 Introduction

It is widely accepted that network security has properties of a public good. A series of works on security games has led to a set of formal tools to analyze the provision of network security by individual agents who control nodes on the network. One distinctive feature of this work over the traditional literature on the provisioning of public goods is the distinction of two types of security technologies, *protection*, which exhibits externalities, and *self-insurance*, which does not. This combination frames network security as a hybrid between a public and a private good, modulated by the relative costs of the two security technologies.

Another distinctive feature of network security over other public goods problems is the existence of uncertainty. An agent's security investment at present only pays off if an attack occurs in a future state of the world. In the existing body of literature, this uncertainty is treated by considering the expected loss as decision variable, largely for the sake of tractability. By contrast, decision science has a rich variety of more realistic models of human and organizational

decision making under uncertainty. One key concept is the notion of risk aversion, typically expressed in a concave utility function. Introducing risk aversion into security games and revisiting the equilibrium strategies is interesting of its own. Even more so because risk aversion naturally leads to a third strategic option, namely agents seeking *market insurance* as a means to transfer the financial risk of uncertain future outcomes.

The question of market insurance for network security has attracted the attention of practitioners, policy makers, and researchers who contributed to a meanwhile sizable body of literature on cyber-insurance.

This paper, to the best of our knowledge for the first time, tries to merge the two streams of research and formally analyzes security games with optional market insurance. To do so, we extend the basic setup of security games with complete information—except for the uncertainty of future states—by a utility function with risk aversion. We discuss this case as an intermediate result before we advance to the analysis of market insurance. The analysis here focuses on the existence of an insurance market, an unanswered question that is relevant to inform policy makers on endeavors to bootstrap a market for cyber-insurance for their assumed positive effects on other frictions to network security not captured by our present formal model, such as information asymmetries and negligence.

This paper is organized as follows: Section 2 recalls the broader context of security games and cyber-insurance within the field of economics of information security. Section 3 presents our model, which is then analyzed in Section 4. The final Section 5 wraps up with discussion and conclusion.

2 Background

An increasing amount of evidence about the economic and technical underpinnings of cybercrime highlight the need for thorough security measures. A growing number of specialized measurement studies demonstrate the professionalism of miscreants concerning a variety of nefarious business models. For example, Holz *et al.* [14] document the elicited trade with payment credentials that have been previously stolen through keylogging malware. Even relatively benign activities such as spam distribution now depend on sophisticated infrastructures existing in the form of botnets and their command-and-control centers [15,21].

The devastating success of these threats frequently depends on interdependencies in computer networks that inhibit the deployment of effective countermeasures. For example, botnets as vehicles behind almost all volume crime on the internet can only exist because some nodes connected to the network apply lower-than-optimal security standards. Similarly, for targeted attacks, a single breach of a corporate perimeter may allow an attacker to harvest resources from all machines located within its confines.

To better understand the implications of these interdependencies for individual defenders, Varian [22] conducted an analysis of system reliability within a public goods game-theoretical framework. He discusses the best effort, weakest-link, and total effort games, as originally analyzed by Hirshleifer [12]. In Varian's

model, security investments take the form of protection effort (e.g., patching system vulnerabilities) where aggregate investments have a decreasing marginal contribution to network security. Further, the security of an individual defender depends on her own effort and on the contributions by all her peers.

Grossklags *et al.* extend this framework by treating security as a hybrid between public and private goods. This is highly interesting because the characteristic as a public or private good is not only determined by the available technology (i.e., its cost) and the architecture of the network (i.e., the functional form of interdependencies). Moreover, individual agents decide strategically on how to split their security investments between protection and self-insurance [7,8]. Self-insurance only affects the investing defender directly, and is consequently a private good (e.g., having good backups). This is different to Ehrlich and Becker's [6] terminology, who use self-insurance to denote loss protection, i.e., reduction of the size of the loss, and protection to denote loss prevention, i.e., reduction of the probability of loss, without differentiating between characteristics of private and public goods. In a more general setting without the distinction between private and public components, both reduce to shifting probability mass in the loss distribution function.¹ Unlike Varian [22], Grossklags *et al.* [7] assume both investment variables to have constant marginal impact across the range of investment opportunities (subject to interdependencies).

Computer security research has been effective in contributing to a better understanding of the uncertainties resulting from attackers' actions. However, this progress in measuring relevant parameters (e.g., attacker intent and attack probabilities) is only partially helpful to understand responders' actions. In particular, we need to have a better grasp of how these factors are perceived by defenders and translated into investment decisions. From behavioral research, it is well-understood that individuals exhibit different risk-coping mechanisms that may depend on a variety of factors (e.g., the amount at stake). Unfortunately, it is rarely the case that risk perception and resulting actions are perfectly in congruence (i.e., risk neutrality).

In fact, for a wide variety of risk scenarios individuals' actions demonstrate risk aversion [11]. Under this behavioral assumption and in the presence of uncertainty, the expected utility of wealth is less than the utility of expected wealth, where the expectations are taken over all possible outcomes of the random future state. To the best of our knowledge there exists no previous work that studies risk-averse agents' decision making in the presence of multiple security investment options (i.e., protection and self-insurance).

In the absence of regulation, institutional behavior is typically more aligned with risk-neutral decision-making, whereas individual decision makers' actions are typically consistent with risk-aversion. Risk aversion and contracts are the

¹ The term self-insurance in the sense of loss protection has also been used by Böhme and Kataria [4] in the context of cyber-insurance describing the option of a single decision maker who operates a large number of computing resources to achieve risk balancing within its own pool of resources rather than joining a risk pool on the insurance market.

only prerequisites for market insurance. More specifically, insurers are offering contracts to risk-averse agents, the insureds.² This risk-pooling should decrease the variance of losses and thereby increase overall welfare [16].

In practice, a number of obstacles have prevented the market for cyber-insurance from achieving maturity. Absence of reliable actuarial data to compute insurance premiums, lack of awareness among decision-makers contributing to too little demand, as well as legal and procedural hurdles have been identified in the “first generation” of cyber-insurance literature until about 2005 [3]. The latter aspect may cause frustration when claiming compensation for damages. Further, entities considering insurance must undergo a series of often invasive security evaluation procedures, revealing their IT infrastructures and policies [1,9]. Meanwhile, witnessing thousands of vulnerabilities, millions of attacks, and substantial improvement in defining security standards and computer forensics calls into question the validity of these factors to causally explain the lack of an insurance market. Consequently, a “second generation” of cyber-insurance literature emerged. Its authors link the market failure with fundamental properties of information technology, specifically correlated risk [2], information asymmetries between insurers and insureds [20], and interdependencies [18,20]. So far, these obstacles have been studied independent of the hybrid private–public good characteristic of network security. Our contribution in this paper is to marry both streams of research and characterize equilibria in a basic model of a security game with market insurance. To keep things tractable, we do not consider correlated risk and we remain in a regime of complete information except for the realization of future losses—see Böhme and Schwartz [5] for a discussion of the validity and implications of these conventions.

3 Model

We devise a stylized game-theoretic model with the intention to focus on the analysis of symmetric equilibria. Occam’s razor was adjusted to emphasize the introduction of risk aversion and the option to obtain market insurance at endogenous but fair premiums. To that end, defenders act as players, attackers as nature, and insurers as mechanism, i. e., price-takers with perfect information about the players’ actions.

3.1 Baseline Security Game

The baseline game includes neither risk aversion nor market insurance. Formally, the base model from which we develop our security games has the following payoff structure. Each of $N \in \mathbb{N}$ players has an initial wealth M_0 . If a given player is

² There are situations where the purchase of insurance might serve as a strategic tool to achieve another purpose. For example, insureds can more credibly threaten with risky behaviors [17]. The purchase of insurance might also help to quell a stakeholder’s fear, uncertainty, and doubt after a security breach. See, for example, banks’ offers of identity theft insurance plans with a free trial period after large-scale data thefts.

attacked and compromised successfully she faces a loss L . Attacks arrive with an exogenous probability of p ($0 \leq p \leq 1$). Players have two security actions at their disposition. Player i chooses a protection level $0 \leq e_i \leq 1$ and a self-insurance level $0 \leq s_i \leq 1$. Finally, $b \geq 0$ and $c \geq 0$ denote the unit cost of protection and self-insurance, respectively.

The post-event wealth function has the following structure:

$$M_1(s_i, e_i; b, c, M_0) = M_0(1 - q \cdot L \cdot (1 - s_i) - be_i - cs_i) \quad (1)$$

where $q \in \{0, 1\}$ is the realization of a random variable indicating loss ($q = 1$) and no loss ($q = 0$). The probability of loss is endogenous and depends on the probability of attack p scaled by the protection effort $H(e_i, e_{-i})$. $H : \mathbb{R}^N \mapsto [0, 1]$ is a contribution function aggregating the protection efforts of player i and all other players (denoted by suffix $-i$). H is monotonically increasing in all its parameters, thereby ensuring that protection generates positive externalities. For the analysis in this paper, we focus on the restricted case in which H describes a weakest link externality, i.e. $H(e_i, e_{-i}) = \min\{e_1, \dots, e_N\}$.

The final utility is mapped to the utility domain by $u = U(M_1)$. As the players maximize expected utility, the combined payoff function of the baseline security game is

$$\begin{aligned} E(u_i) &= p(1 - H(e_i, e_{-i})) \cdot U(M_0(1 - L \cdot (1 - s_i) - be_i - cs_i)) \\ &\quad + (1 - p(1 - H(e_i, e_{-i}))) \cdot U(M_0(1 - be_i - cs_i)). \end{aligned} \quad (2)$$

Post-event wealth is divided into two cases depending on whether a loss occurs or not. In the bad case, new wealth is $M_1 = M_0(1 - L \cdot (1 - s_i) - be_i - cs_i)$. In the good case, new wealth is $M_1 = M_0(1 - be_i - cs_i)$.

3.2 Risk Aversion

Risk aversion is introduced by transforming wealth M_1 to utility $U(M_1)$ using a concave function of type CRRA³,

$$U(M) = \begin{cases} \frac{M^{1-\sigma}}{1-\sigma} & \text{if } \sigma > 0, \sigma \neq 1 \\ \log(M) & \text{if } \sigma = 1, \end{cases} \quad (3)$$

so that $U'(x) = x^{-\sigma}$. $\sigma > 0$ is the degree of risk aversion, an exogenous parameter fixed to $\sigma = 1$ unless otherwise stated. The choice of the CRRA type is convenient because it allows us to derive conclusions that are independent of the initial wealth. This choice also follows established conventions in the cyber-insurance literature (e.g., [2,20]), although CARA-type⁴ utility functions can be found as well [18]. Other researchers are agnostic about the shape of the utility function and just require concavity and twice differentiability [13].

³ CRRA = constant relative risk aversion [19].

⁴ CARA = constant absolute risk aversion.

3.3 Market Insurance

By augmenting the baseline security game with optional market insurance, players will receive an insurance payment, $0 \leq x_i \leq 1$, when a security compromise occurs and they have previously purchased insurance. We assume that agents cannot be overcompensated for losses through a combination of self-insurance and market insurance, i. e., $x_i + s_i \leq 1$. This reflects the principle of indemnity prevalent in the insurance industry. The cost of market insurance, π , is perfectly related to the loss probability and the potential loss in a market with a risk-neutral non-profit insurer who manages a pool of infinitely many homogeneous and independent risks, $\pi = Lp \cdot (1 - H(e_i, e_{-i}))$. However, every realistic (for-profit) insurer would require $\pi > Lp \cdot (1 - H(e_i, e_{-i}))$.

In the presence of market insurance Equation 2 becomes:

$$E(u_i) = p(1 - H(e_i, e_{-i})) \cdot U(M_0(1 - L \cdot (1 - s_i) - be_i - cs_i + x_i(1 - \pi))) + (1 - p(1 - H(e_i, e_{-i}))) \cdot U(M_0(1 - be_i - cs_i - \pi x_i)). \quad (4)$$

Now, in the bad case, new wealth is $M_1 = M_0(1 - L \cdot (1 - s_i) - be_i - cs_i - \pi x_i + x_i)$. In the good case, new wealth is $M_1 = M_0(1 - be_i - cs_i - \pi x_i)$.

3.4 Simplifications

To keep the number of parameters manageable, we assume that $b, c \leq L = 1$, and that, since decisions made on the basis of a CRRA utility function are invariant under multiplicative factors, M_0 can be eliminated. Table 1 in the appendix summarizes all symbols used in our model.

3.5 Payoff Dominance

Theorem 1. *$E[u_i]$ is bounded above by $\max\{1 - b, 1 - c, 1 - \pi, 1 - p\}$. Furthermore, the dominance is strict unless $e_i \in \{0, 1\}$.*

The theorem relies only on the affine structure of our wealth function, together with U being increasing and concave up; a full proof is in the appendix. We use this theorem to help isolate Nash equilibria. If the payoff of each player in a homogeneous strategy achieves the maximizing bound from the theorem, we may conclude that the strategy configuration is a Nash equilibrium.

Conversely, if a strategy configuration results in a utility for some player not conforming to one of the outcomes from the theorem, the only way this configuration can be an equilibrium is if at least one of the outcomes from the theorem is not possible to achieve. This observation can be strengthened by the following corollary.

Corollary 1. *In any hybrid equilibrium where there is a non-zero partial protection investment, the utility of each player is strictly less than $\max\{U(1 - p), U(1 - b), U(1 - c), U(1 - \pi)\}$.*

A proof of the corollary is also in the appendix.

4 Analysis

4.1 Base Model

We begin by briefly reviewing the equilibrium results from the base model (see [10]).

1. Protection equilibria

If $b < p$ and $b < c$, then $(e_i, s_i) = (e_0, 0)$ (protection at level e_0) is a symmetric Nash equilibrium for any e_0 between $\frac{p-c}{p-b}$ and 1.

2. Self-insurance equilibria

If $c < p$ then $(e_i, s_i) = (0, 1)$ (full self-insurance) is a symmetric Nash equilibrium.

3. Passivity equilibria

If $p < c$, then $(e_i, s_i) = (0, 0)$ (passivity) is a symmetric Nash equilibrium.

The above are the only symmetric Nash equilibria for this game. Note that with the exception of partial protection equilibria, all equilibrium strategies are corner strategies. Among all protection equilibrium strategies, the strategy in which each player invests in full protection is Pareto-dominant.

4.2 Base Model with Risk Aversion

Incorporating risk aversion into the base model induces some changes. When the risk-aversion is positive, players have a strong aversion to very low wealth. In fact, for risk aversion coefficients $\sigma \geq 1$, the prospect of having zero wealth results in an infinitely negative utility.⁵ The consequence is that players are no longer satisfied with any strategy in which there is the remote chance of obtaining a non-positive wealth.

We find four distinct types of symmetric Nash equilibrium in the base model supplemented by risk aversion, with $\sigma = 1$.

1. Full protection equilibria

If $b < p$ and $b < c$, then $(e_i, s_i) = (1, 0)$ (full protection) is a symmetric Nash equilibrium.

2. Full self-insurance equilibria

If $c \leq p$ then $(e_i, s_i) = (0, 1)$ (full self-insurance) is a symmetric Nash equilibrium.

3. Partial self-insurance equilibria

If $p < c$, then $(e_i, s_i) = (0, \frac{p}{c})$ (partial self-insurance at the indicated level) is a symmetric Nash equilibrium.

⁵ For $0 < \sigma < 1$ players' utility at zero wealth is finite, but the derivative of utility tends to infinity as wealth approaches zero, so players still have an infinite aversion to retaining a non-positive wealth.

4. Combined protection and self-insurance equilibria

If $p \leq c$, there exists a sufficiently small b such that for any choice of $e_0 < \frac{1-c}{b}$, $(e_i, s_i) = \left(e_0, \frac{p(1-e_0)}{c} + \frac{be_0(c-p(1-e_0))}{c(1-c)} \right)$ (partial protection with partial self-insurance) is a symmetric Nash equilibrium.

An algebraic expression for the maximum b to make this work is difficult to produce (and in fact may not exist), but the existence of b itself follows from the utility function U being differentiable on positive inputs. In the resulting hybrid equilibrium, every player would prefer to invest in full protection because the cost is cheap, but due to the interdependencies inherent in the weakest link game, the maximum investment in protection cannot be set unilaterally, so players are forced to make up for the resulting probability of loss by obtaining self-insurance. If the incentives are such that full self-insurance is desirable, then the incentive to protect will not remain. But if incentives are such that only partial self-insurance is desirable (and if b is sufficiently small), then the configuration with both types of investments is an equilibrium.

Note that passivity is never an equilibrium for any $\sigma \geq 0$, because in such cases players have an infinitely-strong aversion to any non-zero chance of retaining zero wealth.

4.3 Base Model with Risk Aversion and Market Insurance

When we incorporate market insurance, we arrive at more changes. The existence of market insurance ensures that no partial self-insurance investment is optimal. Such partial investments were only possible in the event $p < c$. But if $p < c$ and market insurance is available, then market insurance is always preferable to self-insurance. Even if the reverse inequality holds, it is possible for market insurance to be preferable to self-insurance if there is also a partial protection investment.

1. Full market insurance

If $p \leq c$ then $(e_i, s_i, x_i) = (0, 0, 1)$ (full market insurance) is a symmetric Nash equilibrium.

2. Full self-insurance

If $c \leq p$, then $(e_i, s_i, x_i) = (0, 1, 0)$ (full self-insurance) is a symmetric Nash equilibrium.

3. Full protection

If $b \leq \min\{c, p\}$, then $(e_i, s_i, x_i) = (1, 0, 0)$ (full protection) is a symmetric Nash equilibrium.

4. Partial market insurance and partial self-insurance

If $c = p$, then for any x_0, s_0 with $s_0 + x_0 = 1$, $(e_i, s_i, x_i) = (0, s_0, x_0)$ is a symmetric Nash equilibrium.

5. Partial protection and full market insurance

If $b \leq p$ and $be_0 + p(1 - e_0) < c$, then $(e_i, s_i, x_i) = (e_0, 0, 1)$ is a symmetric Nash equilibrium.




Fig. 1. Overview of feasible symmetric equilibria and corresponding conditions

The last case illustrates an instance in which the availability of market insurance has a positive effect on protection investment. In the same parameter configuration without availability of market insurance, individuals would instead be forced to turn to self-insurance to mitigate against the existing risk. If the additional (compatible) condition $c < p$ is added, then the incentive structure is such that players would prefer to defect to a full self-insurance strategy, and neglect any protection investment.

Figure 1 shows the equilibrium conditions for the case with risk aversion and market insurance.

5 Discussion

In the base model, we find that agents can only with difficulty coordinate on an equilibrium with full protection effort. In particular, the availability of alternative prevention equilibria at $e_0 < 1$ may function to disincentivize defenders to have faith in successful collective preventive actions. As a result, mitigation in the form of full self-insurance may appear more appealing. As risk-neutral decision makers, the agents refrain from security investments when the costs exceed potential losses (see passivity region in Figure 2.a).

Introducing risk aversion for the defender population serves to eliminate the inefficient partial protection equilibria. Further, complete inaction in the form of passivity equilibria disappears. Risk-averse decision makers are willing to invest in security measures costing more than expected losses (see equilibrium strategies for values larger than $pL = p = 0.5$ in Figure 2.b). For example, agents may select a partial self-insurance investment at a fixed level (i. e., $\frac{p}{c}$) when the cost of self-insurance exceeds expected losses.

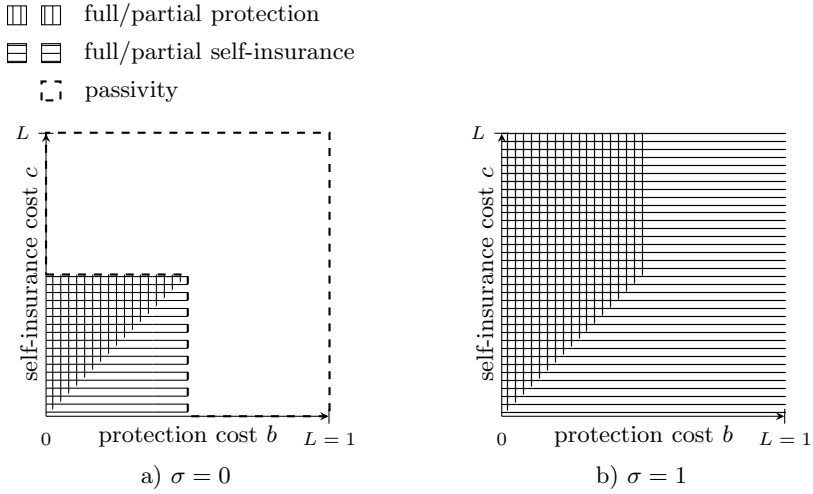


Fig. 2. Symmetric equilibria in the (b, c) -plane for probabilities of attack $p = 0.5$ *without* (left) and *with* (right) risk aversion; no market insurance




Fig. 3. Symmetric equilibria in the (b, c) -plane for different probabilities of attack *with* risk aversion *and* market insurance

In contrast to the base model, we find that equilibria with a joint investment in protection and self-insurance may exist. These outcomes are a more adequate description of reality where a joint defense consisting of prevention and mitigation is common.

The equilibrium conditions including the market insurance option are depicted in Figures 3.a and 3.b. The presence of this third defense strategy serves to clarify the boundaries between the three different defense options. That is, for the most part specific parameter values directly dictate the optimal strategy. Full market insurance, full self-insurance and full protection split the parameter space. However, we observe a hybrid strategy with complementary full market insurance and partial protection investments competing with the full protection equilibrium. Our analysis finds that the hybrid option is payoff-inferior, but might nevertheless be chosen for managerial reasons or inherent unpredictabilities of protection options. Otherwise, full market insurance should only be selected when it is cheaper than both alternative options.

On a more abstract level, our analysis of security games with market insurance can be summarized in three key observations. First, market insurance equilibria exist, and all of them involve full insurance coverage. Second, market insurance is more prevalent for risks with small probability of occurrence. Third, (full) market insurance is a substitute for (expensive) self-insurance technologies, but complementary to (partial and cheap) protection mechanisms.

This leads us to the discussion of limitations of our model and possible extensions. The observation that market insurance responds in a complex manner to the relative cost of protection and self-insurance suggests further investigations are fruitful to account for non-linear cost functions. I.e., protection and self-insurance are likely to exhibit decreasing marginal returns in several scenarios — unlike market insurance which scales linearly as long as the risk is small (and uncorrelated) relative to the pool. The combined equilibrium of partial protection and full market insurance depends on the assumption that the insurer has perfect information about the insureds' protection efforts. If this assumption is relaxed, the arguments made by Shetty *et al.* for the case without self-insurance must be adapted to our security game [20].

Further investigations are needed for the case when insurers charge a strictly positive markup. This will introduce a “gap” of partial market and self-insurance and push the region with full market insurance further to the upper-right corner in Fig. 3. Strictly positive markups are more realistic for various reasons: insurance markets are not fully competitive, regulation requires insurers to be risk-averse, and network risks are often not only independent but correlated as well [5]. Risk correlation leads to longer right tails in the cumulative loss distribution and requires risk-averse insurers to set aside additional safety capital. The cost of this capital has to be added to the fair insurance premium.

To sum up, this paper closes a research gap by modeling network security investments that account for the choice between the hybrid goods of collective protection and individual mitigation and externally provided market insurance. To this end, we have characterized the equilibria of security games with risk

aversion, and security games with risk aversion and market insurance. Overall, as several equilibria with full market insurance exist, market insurance has a place in security games. Moreover, it seems that the missing market problem for cyber-insurance is at least not exacerbated if the agents have a choice between protection and self-insurance.

References

1. Bandyopadhyay, T., Mookerjee, V., Rao, R.: Why IT managers don't go for cyber-insurance products. *Communications of the ACM* 52(11), 68–73 (2009)
2. Böhme, R.: Cyber-insurance revisited. In: *Workshop on the Economics of Information Security (WEIS)*, Cambridge, MA (2005)
3. Böhme, R.: Towards insurable network architectures. *it - Information Technology* 52(5), 290–293 (2010)
4. Böhme, R., Kataria, G.: Models and measures for correlation in cyber-insurance. In: *Workshop on the Economics of Information Security (WEIS)*. University of Cambridge, UK (2006)
5. Böhme, R., Schwartz, G.: Modeling cyber-insurance: Towards a unifying framework. In: *Workshop on the Economics of Information Security (WEIS)*. Harvard University, Cambridge (2010)
6. Ehrlich, I., Becker, G.S.: Market insurance, self-insurance, and self-protection. *Journal of Political Economy* 80(4), 623–648 (1972)
7. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: *Proceedings of the 2008 World Wide Web Conference (WWW 2008)*, Beijing, China, pp. 209–218 (April 2008)
8. Grossklags, J., Christin, N., Chuang, J.: Security and insurance management in networks with heterogeneous agents. In: *Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008)*, Chicago, IL, pp. 160–169 (July 2008)
9. Grossklags, J., Radosavac, S., Cárdenas, A.A., Chuang, J.: Nudge: Intermediaries' Role in Interdependent Network Security. In: Acquisti, A., Smith, S.W., Sadeghi, A.-R. (eds.) *TRUST 2010*. LNCS, vol. 6101, pp. 323–336. Springer, Heidelberg (2010)
10. Grossklags, J.: *Secure or Insure: An Economic Analysis of Security Interdependence and Investment Types*. PhD thesis, University of California, Berkeley (2009)
11. Halek, M., Eisenhauer, J.: Demography of risk aversion. *The Journal of Risk and Insurance* 68(1), 1–24 (2001)
12. Hirshleifer, J.: From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice* 41(3), 371–386 (1983)
13. Hofmann, A.: Internalizing externalities of loss prevention through insurance monopoly: An analysis of interdependent risks. *Geneva Risk and Insurance Review* 32(1), 91–111 (2007)
14. Holz, T., Engelberth, M., Freiling, F.: Learning More About the Underground Economy: A Case-Study of Keyloggers and Dropzones. In: Backes, M., Ning, P. (eds.) *ESORICS 2009*. LNCS, vol. 5789, pp. 1–18. Springer, Heidelberg (2009)
15. Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., Savage, S.: Spamalytics: An empirical analysis of spam marketing conversion. In: *Proceedings of the Conference on Computer and Communications Security (CCS)*, Alexandria, VA (October 2008)

16. Kesan, J., Majuca, R., Yurcik, W.: The economic case for cyberinsurance. In: Proceedings of the Fourth Workshop on the Economics of Information Security (WEIS), Cambridge, MA (June 2005)
17. Kirstein, R.: Risk neutrality and strategic insurance. The Geneva Papers on Risk and Insurance 25, 251–261 (2000)
18. Ogut, H., Menon, N., Raghunathan, S.: Cyber insurance and IT security investment: Impact of interdependent risk. In: Fourth Workshop on the Economics of Information Security (WEIS), Cambridge, MA (June 2005)
19. Pratt, J.: Risk aversion in the small and in the large. *Econometrica* 32(1-2), 122–136 (1964)
20. Shetty, N., Schwartz, G., Felegyhazi, M., Walrand, J.: Competitive Cyber-Insurance and Internet Security. In: Workshop on Economics of Information Security 2009. University College London, England (2009)
21. Stone-Gross, B., Holz, T., Stringhini, G., Vigna, G.: The underground economy of spam: A botmaster’s perspective of coordinating large-scale spam campaigns. In: Proceedings of the 4th USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), Boston, MA (March 2011)
22. Varian, H.: System reliability and free riding. In: Camp, J., Lewis, S. (eds.) *Economics of Information Security*. *Advances in Information Security*, vol. 12, pp. 1–15. Kluwer Academic Publishers, Dordrecht (2004)

Appendix

Table 1. List of Symbols

| Symbol | Type | Meaning | Constraints |
|----------|-----------------|---|--------------------|
| b | parameter | cost of protection | $0 < b \leq 1$ |
| c | parameter | cost of self-insurance | $0 < c \leq 1$ |
| e_i | choice variable | level of player i 's protection | |
| E | operator | expected value (over loss realization) | |
| H | function | protection contribution function | |
| L | constant | size of the loss | $L = 1$ |
| M_0 | constant | initial wealth | eliminated |
| M_1 | variable | ex-post wealth | |
| N | parameter | number of players | $N > 1$ |
| p | parameter | probability of loss | |
| π | variable | cost of market insurance | |
| q | random variable | realization of the loss | $q \in \{0, 1\}$ |
| s_i | choice variable | level of player i 's self-insurance | $s_i + x_i \leq 1$ |
| σ | parameter | risk aversion | $\sigma \geq 0$ |
| u_i | variable | player i 's utility | |
| U | function | utility function | |
| x_i | choice variable | level of player i 's market insurance | $s_i + x_i \leq 1$ |

A Proof of Theorem 1

Proof. Assume that player strategies comprise a symmetric Nash equilibrium. Let e, s, x be the homogeneous protection, self-insurance, and market insurance investments, respectively. Then, we can write the expected utility of player i as

$$\begin{aligned}
 E[u_i] &= p(1-e) \cdot U(s+x-be-cs-\pi x) + (1-p(1-e)) \cdot U(1-be-cs-\pi x) \\
 &\leq U(p(1-e) \cdot (s+x-be-cs-\pi x)) + (1-p(1-e)) \cdot (1-be-cs-\pi x) \\
 &= U(p(1-e)(s+x) + p(1-e)(-be-cs-\pi x)) \\
 &\quad + (1-p(1-e) + (1-p(1-e))(-be-cs-\pi x)) \\
 &= U(p(1-e)(s+x) + (-be-cs-\pi x) + 1-p(1-e)) \\
 &= U(ps+px-pes-pex-be-cs-\pi x+1-p+pe) \\
 &= U(1-p+e(p-b) + x(p-\pi) + s(p-c) - ep(s+x)).
 \end{aligned}$$

Since U is increasing we can maximize the last formula in the derivation above by choosing e, s, x to maximize the quantity inside the U function.

Excluding the last term, that formula is linear; and the last term is strictly negative whenever at least one of e or $s+x$ is positive. So the choice if e, s, x to maximize the formula can be easily determined from $\min\{p, b, c, \pi\}$ – namely, we choose $(e, s, x) = (0, 0, 0)$ if p is the minimum, resulting in utility $U(1-p)$; we choose $(e, s, x) = (1, 0, 0)$ if b is the smallest, obtaining utility $U(1-b)$; we choose $(e, s, x) = (0, 1, 0)$ if c is the least obtaining utility $U(1-c)$; and if π is the min we choose $(e, s, x) = (0, 0, 1)$, obtaining utility $U(1-\pi)$. If there are equalities among terms, then the proper choice of e, s, x to maximize the formula is not uniquely determined, but there is nothing about equality that would change the final utility. We conclude that for any choice of e, s, x , the expected utility of each player $E[u_i]$ cannot exceed $\max\{U(1-p), U(1-b), U(1-c), U(1-\pi)\}$.

For the strictness result, observe that the inequality in the second step follows from the fact that U is concave down. The only time that inequality is an equality is when one of the scaling factors $p(1-e)$ or $1-p(1-e)$ is zero. Since we have assumed $p > 0$, the only way to have equality is if $e \in \{0, 1\}$.

B Proof of Corollary 1

Proof. For the corollary, we first note that in any hybrid equilibrium in which there is a partial protection investment, we must necessarily have $\min\{b, p\} < \min\{c, \pi\}$. Otherwise, any player could unilaterally make an investment in full self-insurance or full market insurance and achieve the maximum bound from the theorem, which would necessarily be an improvement due to the strict inequality. Hence one of p or b minimizes $\{p, b, c, \pi\}$. If $p \leq b$, then setting $e > 1$ results in the final term in the derivation above being less than $U(1-p)$. On the other hand, if $b \leq p$, then any investment $e < 1$ results in a final term in the above derivation being strictly less than $U(1-b)$. [Investment in s or x (assuming $s+x \leq 1$) cannot make up for this, because the last term in the last formula subtracts the advantage gained from the protection investment].