

Secure or Insure?

A Game-Theoretic Analysis of Information Security Games

Jens Grossklags
UC Berkeley
School of Information
Berkeley, CA 94720
jensg@ischool.berkeley.edu

Nicolas Christin
Carnegie Mellon University
INI/CyLab Japan
Kobe, 650-0044 Japan
nicolasc@cmu.edu

John Chuang
UC Berkeley
School of Information
Berkeley, CA 94720
chuang@ischool.berkeley.edu

ABSTRACT

Despite general awareness of the importance of keeping one's system secure, and widespread availability of consumer security technologies, actual investment in security remains highly variable across the Internet population, allowing attacks such as distributed denial-of-service (DDoS) and spam distribution to continue unabated. By modeling security investment decision-making in established (e.g., weakest-link, best-shot) and novel games (e.g., weakest-target), and allowing expenditures in self-protection versus self-insurance technologies, we can examine how incentives may shift between investment in a public good (protection) and a private good (insurance), subject to factors such as network size, type of attack, loss probability, loss magnitude, and cost of technology. We can also characterize Nash equilibria and social optima for different classes of attacks and defenses. In the weakest-target game, an interesting result is that, for almost all parameter settings, more effort is exerted at Nash equilibrium than at the social optimum. We may attribute this to the "strategic uncertainty" of players seeking to self-protect at just slightly above the lowest protection level.

Categories and Subject Descriptors

C.2 [Computer Systems Organization]: Computer-Communication Networks; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*; K.4.4 [Computers and Society]: Electronic Commerce—*Security*

General Terms

Economics, Reliability, Security

Keywords

Economics of the Internet, Game Theory, Public Goods, Incentive-Centered Design and Engineering, Security, Protection, Self-Insurance

1. INTRODUCTION

The Internet has opened new and attractive channels to publicize and market products, to communicate with friends and colleagues, and to access information from spatially distributed resources. Though it has grown significantly, the network's architecture still reflects the cooperative spirit of its original designers [32]. Unfortunately, today's network users are no longer held together by that same sense of camaraderie and common purpose. For instance, concrete evidence of the tragedy of the commons [21] occurring in

This is a preprint of an article accepted for publication in the 17th International World Wide Web Conference (WWW'08), Beijing, China, April 21–25 2008. Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

peer-to-peer filesharing networks has been documented for a long time [2]. Accordingly, studies of networking protocols and user interaction have been assuming users to be selfish and to act strategically [37].

Selfish users are one thing, but the expansion of the Internet has also attracted individuals and groups with often destructive motivations; these "attackers" intend to improve on their perceived utility by exploiting or creating security weaknesses and harming or inconveniencing other network users. Some malicious entities are motivated by peer recognition, or curiosity, and are often undecided regarding the ethical legitimacy of their behavior [19, 20]. Others have clearly demonstrated financial goals [17]. Problematic behaviors and threats include attacks on the network as a whole, attacks on selected end-points, undesirable forms of interactions such as spam e-mail, and annoyances such as Web pages that are unavailable or defaced. As a result, users cannot rely and trust other network participants [13].

When asked in surveys, network users say they are interested in preventing attacks and mitigating the damages from computer and information security breaches [1, 40]. Researchers and industry have responded by developing numerous security technologies to alleviate many of the aforementioned problems [4], thereby expecting to help improving individual security practices.

Nevertheless, security breaches are common, widespread and highly damaging. The "I Love You" virus [27], Code Red [29] and Slammer worms [28], to cite the most famous cases, have infected hundreds of thousands of machines and caused, all together, billions of dollars in damages. Underground markets for processor time on compromised end-systems are developing [17] thanks to large population of home computers that can be easily commandeered by third-parties. The high financial impact of security failures is explained by user surveys [6, 10], which show strong evidence that comprehensive security precautions, be they patching, spyware-removal tools, or even sound backup strategies, are missing from a vast majority of systems surveyed.

In other words, despite a self-professed interest in security, most individuals do not implement effective security on their systems, even though necessary technologies and methods are (by and large) readily available. We propose to investigate the root causes of the disconnect between users' actions and their intentions.

In practice, there is a large variety of situations in which users face security threats, and an equally large number of possible responses to threats. However, we postulate in this paper that one can model most security interactions through a handful of "security games," and with a small number of decision parameters upon which each user can act.

More precisely, building upon public goods literature [23, 43], we consider the classical best shot, total effort, and weakest-link

games, and will analyze them in a security context. We complement these three games with a novel model, called the “weakest-target” game, which allows us to describe a whole class of attacks ranging from insider threats to very aggressive worms. Furthermore, while most research on the economics of security focuses on security investments as a problem with a single variable (e.g., amount of money spent on security), our analysis is the first to decouple protection investments (e.g., setting up a firewall) from insurance coverage (e.g., archiving data as back up). This decoupling allows us to explain a number of inefficiencies in the observed user behaviors.

This paper is only a first step toward a more comprehensive modeling of user attitudes toward security issues. Indeed, the present study relies on game theory, mostly using Nash equilibrium and social optima concepts. As such, we primarily view this study as a theoretical basis for follow up experimental work using laboratory experiments with human participants. We nevertheless show that the models and results derived here provide for considerable insights.

The rest of this paper is organized as follows. We elaborate in Section 2 the relationship of our work with related research, before introducing our game-theoretic models in Section 3. We present an analysis of the Nash equilibria (Section 4) and social optima (Section 5) for all of these games, and discuss our findings in Section 6. We conclude in Section 7.

2. RELATED WORK

The economics of information security is a growing research area with a diverse set of participating researchers from various disciplines. Important common anchors are the observations that misaligned incentives and positive and negative externalities play significant roles in the strategies used by each party in the battle between attackers and potential victims [4, 5].

Economics as a tool for security analysis has gained in importance since the economy of attackers has become increasingly rational (e.g., motivated by greed), over the last years [17]. This increasingly rational behavior stands in contrast to that exhibited by the hacker communities of the 1980s and 1990s, who valued reputation, intellectual achievement, and even entertainment above financial incentives [19, 20].

Most of the initial results obtained in security economics research concern the analysis of optimal security investments. For example, Gordon and Loeb [18] as well as Hausken [22] focus on the impact of different security breach functions and degrees of vulnerability on an entity’s investment strategy. More specialized models have been proposed to analyze a subset of important security management problems. For instance, August and Tunca scrutinize optimal system update strategies when patching a system against security vulnerabilities is costly [7]. Rescorla investigates the impact of code quality control on vulnerability of software [31]. Böhme and Kataria study individual security investment decisions and market insurance offers when correlation of cyber-risks within a single firm and across multiple firms differ [8].

From a policy standpoint, Bull et al. [11] argue that given the state of heterogeneous networks no single security policy will be applicable to all circumstances. They suggest that, for a system to be viable from a security standpoint, individuals need to be empowered to control their own resources and to make customized security trade-offs. This stands in contrast to the traditional centralized structure where all security decision are made by a central planner (e.g., the IT department). Nevertheless, as Anderson suggests, organizational and structural dependencies have to be considered in individual security decision making [3].

While many models prescribe behavior in individual choice situations, the focus of our work is to model and study strategic interaction with respect to security decisions in networked systems, in an effort to understand the impact of individual choices on a larger group. Such interaction usually involves common as well as conflicting interests. (Pure conflict, in which the interests of the two antagonists are completely opposed, is a special case.) This mutual dependence as well as opposition guarantees for a much richer scenario for analysis [35].

To better understand the implications of this mutual dependence, Varian [43] conducts 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 Hirschleifer [23]. The main difference from classical public goods theory is that within the framework of computer reliability “considerations of costs, benefits, and probability of failure become paramount, with income effects being a secondary concern.” [43] Varian focuses on two-player games with heterogeneous effort costs and benefits from reliability.¹ He also adds an inquiry into the role of taxes and fines, and differences between simultaneous and sequential moves.

Our work generalizes [43] in several aspects. First, instead of considering security decisions to be determined by a single “security” variable, we identify two key components of a security strategy: self-protection (e.g., patching system vulnerabilities) and self-insurance (e.g., having good backups). More precisely, we allow agents to self-protect and/or self-insure their resources in N -player games. We also contrast the three canonical games discussed by Varian with two more complex “weakest-target” games that represent a more complicated incentive structure, which we believe applies to a whole class of security issues.

Outside the information security context, the dual role of self-protection and self-insurance was first recognized by [15]. To provide a more precise definition, self-protection stands for the ability to reduce the probability of a loss – for example, by installing a firewall application which limits the amount of traffic allowed to communicate with one’s network. Self-insurance, on the other hand, denotes a reduction in the magnitude of a loss, e.g., by performing regular backups on existing data. Some technologies and practices such as disconnecting a computer from a network do both. Ehrlich and Becker [15] focus in their analysis on the comparison of self-protection and self-insurance to market insurance. They find that, for rare loss events, there is less incentive to self-insure losses than to use market insurance. This is due to their assumption, that the price of self-insurance is independent of the probability of the loss. An additional result is that the demand for self-insurance grows with the base loss of a security threat. As an outcome of their work, they characterize self-insurance and market insurance as substitutes, and self-protection and market insurance as complements. Our analysis complements the work in [15] by extending the concepts of self-protection and self-insurance to the public goods and security context.

3. DESCRIPTION OF SECURITY GAMES

We define a security game as a game-theoretic model that captures essential characteristics of decision making to protect and self-insure resources within a network. Varian [43] observed that frequently the success of security (or reliability) decision making depends on a joint protection level determined by all participants

¹A distinction between reliability and security, in terms of consequences, may exist [24]. In this study, we do not follow this distinction and consider reliability as a key component of security.

of a network. The computation of the protection level will often take the form of a public goods contribution function with nonrival and nonexcludable benefits or consequences. A main observation is that dependent on the contribution function individuals may be able to freeride on others' efforts. However, individuals may also suffer from inadequate protection efforts by other members if those have a decisive impact on the overall protection level.

Following Varian's exposition, we analyze three canonical contribution functions that determine a global protection level. Different from Varian's work however, here network members have a second action available: They can decide to self-insure themselves from harm. The success of insurance decisions is completely independent of protection choices made by the individual and others. Consequently, the games we consider share qualities of private (on the insurance side) and public (on the protection side) goods. We further add to the research literature by studying two additional games with a more complex determination of protection levels.

Security games share the following key assumptions: (i) all entities in the network share a single purely public protection output, (ii) a single individual decides on protection efforts for each entity (so we do not assume a second layer of organizational decision making), (iii) protection costs per unit are identical for each entity, and (iv) all decisions are made simultaneously. These assumptions are commonly made also in models on decision making of partners in military alliances [33]. We add to these main assumptions that individuals are able to self-insure resources at a homogeneous cost with self-insurance being a purely private good.

Formally, the basic model from which we develop the security games has the following payoff structure. Each of $N \in \mathbb{N}$ players receives an endowment M . If she is 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 an insurance level $0 \leq s_i \leq 1$ and a protection level $0 \leq e_i \leq 1$. Finally, $b \geq 0$ and $c \geq 0$ denote the unit cost of protection and insurance, respectively. The generic utility function has the following structure:

$$U_i = M - pL(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i, \quad (1)$$

where following usual game-theoretic notation, e_{-i} denotes the set of protection levels chosen by players other than i . H is a "contribution" function that characterizes the effect of e_i on U_i , subject to the protection levels chosen (contributed) by *all* other players. We require that H be defined for all values over $(0, 1)^N$. However, we do not place, for now, any further restrictions on the contribution function (e.g., continuity). From Eqn. (1), the magnitude of a loss depends on three factors: i) whether an attack takes place (p), ii) whether the individual invested in self-insurance ($1 - s_i$), and iii) the magnitude of the joint protection level ($1 - H(e_i, e_{-i})$). Self-insurance always lowers the loss that an individual incurs when compromised by an attack. Protection probabilistically determines whether an attack is successful. Eqn. (1) therefore yields an expected utility.

We introduce five games in the following discussion. In selecting and modeling these games we paid attention to comparability of our security games to prior research (e.g., [23, 33, 43]). The first three specifications for H represent important baseline cases recognized in the public goods literature. To allow us to cover most security dilemmas, we add two novel games, for which we could not find a formal representation in the literature. All games are easy to interpret within and outside the online security context.

Total effort security game: The global protection level of the network depends on the sum of contributions normalized over the number of all participants. That is, we define $H(e_i, e_{-i}) = \frac{1}{N} \sum_i e_i$,

so that Eqn. (1) becomes

$$U_i = M - pL(1 - s_i)(1 - \frac{1}{N} \sum_k e_k) - be_i - cs_i. \quad (2)$$

Economists identified the sum of efforts (or total effort) contribution function long before the remaining cases included in this paper [23]. We consider a slight variation of this game to normalize it to the desired parameter range. A typical parable for the total sum function is that the effectiveness of a dam or city wall depends on its strength that is contributed to by all players. In terms of security the average contributions matters if an attacker wants to successfully conquer the majority of machines in a network one-by-one. For instance, consider a building plan for a new technology that is spread across a company's network and which is considerably more valuable to an attacker, if obtained in its entirety.

As another example, maybe more related to Internet security, consider parallelized file transfers, as in the BitTorrent peer-to-peer service. It may be the case that an attacker wants to slow down transfer of a given piece of information; but the transfer speed itself is a function of the aggregate effort of the machines participating in the transfer. Note that, the attacker in that case is merely trying to slow down a transfer, and is not concerned with completely removing the piece of information from the network: censorship actually results in a different, "best shot" game, as we discuss later.

Weakest-link security game: The overall protection level depends on the minimum contribution offered over all entities. That is, we have $H(e_i, e_{-i}) = \min(e_i, e_{-i})$, and Eqn. (1) takes the form:

$$U_i = M - pL(1 - s_i)(1 - \min(e_i, e_{-i})) - be_i - cs_i. \quad (3)$$

This game describes the situation where a levee or city wall that is too low at any point leads to a negative payoff to all players in the event of a flood or attack. The weakest link game is easily the most recognized public goods problem in computer security by business professionals and researchers alike.² Once the perimeter of an organization is breached it is often possible for attackers to leverage this advantage. This initial compromise can be the result of a weak password, an inconsistent security policy, or some malicious code infiltrating a single client computer.

Best shot security game: In this game, the overall protection level depends on the maximum contribution offered over all entities. Hence, we have $H(e_i, e_{-i}) = \max(e_i, e_{-i})$, so that Eqn. (1) becomes

$$U_i = M - pL(1 - s_i)(1 - \max(e_i, e_{-i})) - be_i - cs_i. \quad (4)$$

As an example of a best shot game, consider a set of walls of which the highest sets the effectiveness benchmark. Among information systems, networks with built-in redundancy, such as peer-to-peer, sensor networks, or even Internet backbone routes, share resilience qualities with the best shot security game; for instance, to completely take down communications between two (presumably highly connected) backbone nodes on the Internet, one has to shut down all possible routes between these two nodes. Censorship-resistant networks are another example of best shot games. A piece of information will remain available to the public domain as long

²See, for example, see a recent interview with a security company CEO. New York Times (September 12, 2007), "Who needs hackers," available at <http://www.nytimes.com/2007/09/12/technology/techspecial/12threat.html>. Stating that: "As computer networks are cobbled together [...] the Law of the Weakest Link *always* seems to prevail."

as a single node serving that piece of information can remain unharmed [14].

Weakest-target security game (without mitigation): Here, an attacker will *always* be able to compromise the entity (or entities) with the lowest protection level, but will leave other entities unharmed. This game derives from the security game presented in [12]. Formally, we can describe the game as follows:

$$H(e_i, e_{-i}) = \begin{cases} 0 & \text{if } e_i = \min(e_i, e_{-i}), \\ 1 & \text{otherwise,} \end{cases} \quad (5)$$

which leads to

$$U_i = \begin{cases} M - pL(1 - s_i) - be_i - cs_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M - be_i - cs_i & \text{otherwise.} \end{cases} \quad (6)$$

The weakest-target game markedly differs from the weakest link. There is still a decisive security level that sets the benchmark for all individuals. It is determined by the individual(s) with the lowest chosen effort level. However, in this game all entities with a protection effort strictly larger than the minimum will remain unharmed.

In information security, this game captures the situation in which an attacker is interested in securing access to an arbitrary set of entities with the lowest possible effort. Accordingly, she will select the machines with the lowest security level. An attacker might be interested in such a strategy if the return on attack effort is relatively low, for example, if the attacker uses a compromised machine to distribute spam. Such a strategy is also relevant to an attacker with limited skills, a case getting more and more frequent with the availability of automated attack toolboxes [41]; or, when the attacker's goal is to commandeer the largest number of machines using the smallest investment possible [17]. Likewise, this game can be useful in modeling insider attacks – a disgruntled employee may for instance very easily determine how to maximize the amount of damage to her corporate network while minimizing her effort.

Weakest-target security game (with mitigation): This game is a variation on the above weakest-target game. The difference is that, the probability that the attack on the weakest protected player(s) is successful is now dependent on the security level $\min e_i$ chosen. That is,

$$H(e_i, e_{-i}) = \begin{cases} 1 - e_i & \text{if } e_i = \min(e_i, e_{-i}), \\ 1 & \text{otherwise,} \end{cases} \quad (7)$$

so that

$$U_i = \begin{cases} M - pL(1 - s_i)(1 - e_i) - be_i - cs_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M - be_i - cs_i & \text{otherwise.} \end{cases} \quad (8)$$

This game represents a nuanced version of the weakest-target game. Here, an attacker is not necessarily assured of success. In fact, if all individuals invest in full protection, not a single machine will be compromised. This variation allows us to capture scenarios where, for instance, an attacker targets a specific vulnerability, for which an easily deployable countermeasure exists.

Limitations: With the analysis in this paper we aim for a more thorough understanding of the ecology of security threats and defense functions an individual or organization faces and has to respond to. We have generalized and newly developed models that represent vastly different security scenarios and will call for different actions. As Hirshleifer observed [23] a security practitioner will be presented with “all kinds of intermediate cases and combinations,” e.g., social composition functions involving all of these five rules as well as other not identified yet. Some minor variations would be the “location of the top decile, or the total of the best three

shots, or the average of the best and worst shots, or the variance or skewness” etc. See also [7] and [26] for variations in which the likelihood of a compromise depends on the number of unprotected players.

4. NASH EQUILIBRIUM ANALYSIS

We next determine the equilibrium outcomes where each individual chooses protection effort and self-insurance investments unilaterally, in an effort to maximize her own utility. In Section 5, we then compare these results to the protection efforts and self-insurance levels chosen if coordinated by a social planner.

4.1 Total effort

Let us focus on player i , and consider e_k for $k \neq i$ as exogenous. Then, U_i is a function of two variables, e_i and s_i . From Eqn. (2), U_i is twice differentiable in e_i and s_i , with $\partial^2 U_i / \partial s_i^2 = 0$ and $\partial^2 U_i / \partial e_i^2 = 0$. Hence, according to the second derivative test, only $(e_i, s_i) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$ can be an extremum – that is, possible Nash equilibria are limited to these four values (or to strategies yielding a payoff constant regardless of e_i and/or s_i). As long as at least one of b or c is strictly positive, $(e_i, s_i) = (1, 1)$ is always dominated by either $(e_i, s_i) = (1, 0)$ or $(e_i, s_i) = (0, 1)$ and cannot define a Nash equilibrium. Let us analyze the three other cases:

- $(e_i, s_i) = (0, 0)$. Replacing in Eqn. (2), we get

$$U_i = M - pL \left(1 - \frac{1}{N} \sum_{k \neq i} e_k \right). \quad (9)$$

- $(e_i, s_i) = (0, 1)$. Replacing in Eqn. (2), we get

$$U_i = M - c. \quad (10)$$

- $(e_i, s_i) = (1, 0)$. Replacing in Eqn. (2), we get

$$U_i = M - pL \left(1 - \frac{1}{N} - \frac{1}{N} \sum_{k \neq i} e_k \right) - b. \quad (11)$$

Result 1: After investigating Eqs. (9–11) we can identify three Nash equilibrium strategies.

- **Full protection eq.:** If $pL > bN$ and $c > b + pL \frac{N-1}{N}$, meaning that protection is cheap, potential losses are high, and insurance is extremely overpriced, then the (only) Nash equilibrium is defined by everybody protecting but not insuring, that is, $(e_i, s_i) = (1, 0)$.
- **Full self-insurance eq.:** In the other cases where $pL > bN$, $(e_i, s_i) = (0, 1)$ is a Nash equilibrium. Also, if $c < pL < bN$ (expected losses above insurance costs), then $(e_i, s_i) = (0, 1)$, is a Nash equilibrium.
- **Passivity eq.:** If $pL < bN$ and $pL < c$, then the expected losses are small enough so that complete passivity, defined by $(e_i, s_i) = (0, 0)$ for all players, is a Nash equilibrium.

Increasing number of players N: As the number of players increases, protection equilibria become more and more unlikely to occur. Indeed, in a total effort scenario, “revenues” yielded by a player's investment in security have to be shared with all of the other participants, making it an increasingly uninteresting strategy for the player as the network grows.

4.2 Weakest-link

Let $e_0 = \min_i(e_i)$. From Eqn. (3), we have $U_i = M - pL(1 - s_i)(1 - e_0) - be_i - cs_i$, so that $\frac{\partial U_i}{\partial s_i} = pL(1 - e_0) - c$, and, for all i ,

$$U_i \leq M - pL(1 - s_i)(1 - e_0) - be_0 - cs_i,$$

which is reached for $e_i = e_0$. So, in a Nash equilibrium, everybody picks the same $e_i = e_0$. It follows that Nash equilibria are of the form $(e_0, 0)$ or $(0, 1)$.

Result 2: *In the weakest link security game, we can identify three types of Nash equilibrium strategies. However, there exist multiple pure protection equilibria.*

Denote by \hat{e}_0 the minimum of the protection levels initially chosen by all players. We have

- **Multiple protection equilibria:** If $pL > b$ and $\{(\hat{e}_0 > (pL - c)/(pL - b) \text{ for } c < pL) \cup (pL \geq c)\}$, then $(e_i, s_i) = (\hat{e}_0, 0)$ for all i is a Nash equilibrium: everybody picks the same minimal security level, but no one has any incentive to lower it further down. This equilibrium can only exist for $b \leq c$, and may be inefficient, as it could be in the best interest of all parties to converge to $e_i = 1$, as we discuss later in Section 5.
- **Full self-insurance eq.:** If $pL > c$ and $\{\hat{e}_0 < (pL - c)/(pL - b) \cup b > pL\}$, then $(e_i, s_i) = (0, 1)$ for all i is a Nash equilibrium: essentially, if the system is not initially secured well enough (by having all parties above a fixed level), players prefer to self-insure.
- **Passivity eq.:** If $pL < b$ and $pL < c$, then $(e_i, s_i) = (0, 0)$ is the only Nash equilibrium – both insurance and protection are too expensive.

Notice that if $\hat{e}_0 = (pL - c)/(pL - b)$, then both full self-insurance $((e_i, s_i) = (0, 1)$ for all i) and protection $((e_i, s_i) = (\hat{e}_0, 0)$ for all i) form a Nash equilibrium. In particular, if $b = c$ ($b < pL$ and $c < pL$) full protection $(e_i, s_i) = (1, 0)$ and full self-insurance $(e_i, s_i) = (0, 1)$ are Nash strategies.

Increasing number of players N: The weakest link security game, much like the tacit coordination game of [42] has highly volatile protection equilibria when the number of players increase. In fact, any protection equilibrium has to contend with the strategic certainty of a self-insurance equilibrium. To view this, consider the cumulative distribution function $F(e_i)$ over the protection strategies e_i of a given player i . From what precedes, with pure strategies, in the Pareto-optimum, $F(1) = 1$ and $F(e_i) = 0$ for $e_i < 1$. Assuming all N players use the same c.d.f. F , then the c.d.f. of $e_0 = \min_i\{e_i\}$ is given by $F_{\min}(e_0) = 1 - (1 - F(e_0))^N$ [42]. So, $F_{\min}(1) = 1$ and $F_{\min}(e_0) = 0$ for $e_0 < 1$ as well. Now, assume there is an arbitrarily small probability $\varepsilon > 0$ that one player will defect, that is $F(0) = \varepsilon$. Then, $F_{\min}(0)$ converges quickly to 1 as N grows large. That is, it only takes the slightest rumor that one player may defect for the whole game to collapse to the $(e_i, s_i) = (0, 1)$ equilibrium.

4.3 Best shot

Let $e^* = \max_i(e_i)$. Eqn. (4) gives

$$U_i = M - pL(1 - s_i)(1 - e^*) - be_i - cs_i.$$

Clearly, $(e_i, s_i) = (1, 1)$ is suboptimal, so that three strategies may yield the highest payoff to user i .

- Selecting $(e_i, s_i) = (0, 0)$ yields $U_i = M - pL(1 - e^*)$.

- Selecting $(e_i, s_i) = (1, 0)$ yields $U_i = M - b$.
- Selecting $(e_i, s_i) = (0, 1)$ yields $U_i = M - c$.

Result 3: *From the above relationships, we can identify the following pure Nash equilibrium strategies.*

- **Full self-insurance eq.:** If $b > c$ we find that the self-insurance equilibrium $(\forall i, (e_i, s_i) = (0, 1))$ is the only possible Nash equilibrium.
- **Passivity eq.:** If $pL < b$ and $pL < c$ agents prefer to abstain from security actions $(\forall i, (e_i, s_i) = (0, 0))$.

In particular, there is no protection equilibrium in this game. For one protection equilibrium to exist, we would need $b < c$ and $pL > b$. But even assuming that this is the case, as long as the game is synchronized, players endlessly oscillate between securing as much as possible ($e_i = 1$) and free-riding ($e_i = 0$). This is due to the fact that as soon as one player secures, all others have an incentive to free-ride. Conversely, if everybody free-rides, all players have an incentive to deviate and secure as much as possible.

Increasing number of players N: In the absence of coordination between players, the outcome of this game is globally independent of the number of players N , as there is no protection equilibrium, and the insurance equilibrium is independent of the number of players. However, the game may be stabilized by using player coordination (e.g., side payments) for low values of N , something harder to do as N grows.

4.4 Weakest-target (without mitigation)

Fix the strategy point and let $\varepsilon < \frac{pL}{2b}$. Let e_0 be the minimum effort level of any player. Then no player selects a higher effort than $e_0 + \varepsilon$ because it dominates all higher effort levels. However, any player at e_0 would prefer to switch to $e_0 + 2\varepsilon$. Then the change in her payoff is greater than $pL - 2\frac{pL}{2b}b = 0$. Because this deviation is profitable this strategy point is not an equilibrium.³

Result 4: *In the weakest-target game with an attacker of infinite strength we find that pure Nash equilibria for non trivial values of b, p, L and c do not exist.*

Mixed strategy equilibria. While no pure Nash equilibria exist, let us explore the existence of a mixed strategy equilibrium. We use the shorthand notation $e_i = e, s_i = s$ here, and consider mixed strategies for choosing e . There are two cases to consider.

Case $c > pL$: If $c > pL$ then dominance arguments immediately lead to $s = 0$ meaning that nobody buys any self-insurance.

An equilibrium strategy may be parametrized by e . For a given player, the utility function U becomes a function of a single variable e . Let $f(e)$ be the probability distribution function of effort in the weakest-target game and let $F(e)$ be the cumulative distribution function of effort. Assuming only one player is at the minimum protection level, shall an attack occur, the probability of being the victim is then $(1 - F(e))^{N-1}$. (All N players choose protection levels greater than e .)

Then the utility is given by

$$U = M - pL(1 - F(e))^{N-1} - be. \quad (12)$$

³While this proof assumes the player is initially at $(e_i, s_i) = (e_0, 0)$, it can be trivially extended to the case $(e_i, s_i) = (e_0, s)$ with $s > 0$ by picking $\varepsilon < \frac{pL}{2b}(1 - s) + \frac{cs}{2b}$ for any s in $(0, 1]$.

In a Nash equilibrium, the first-order condition $dU/de = 0$ must hold, so that:

$$(N-1)pLf(e)[1-F(e)]^{N-2} - b = 0$$

If we substitute $G = (1 - F(e))$ and $g = -f$ we can write $G^{N-2}dG/de = -b/p(N-1)L$, which, by integration yields

$$\int_{G(e)}^{G(0)} G^{N-2}dG = \int_e^0 \frac{-b}{p(N-1)L} d\hat{e},$$

that is

$$G^{N-1} \Big|_{G(e)}^{G(0)} = \frac{-b}{pL} e. \quad (13)$$

With $G(0) = 1$,

$$G(e) = \left(1 - \frac{b}{pL}e\right)^{\frac{1}{N-1}}.$$

Differentiating, we get

$$g(e) = -\frac{1}{N-1} \frac{b}{pL} \left(1 - \frac{b}{pL}e\right)^{-\frac{N-2}{N-1}},$$

and, replacing $g = -f$ we find,

$$f(e) = \frac{1}{N-1} \frac{b}{pL} \left(1 - \frac{b}{pL}e\right)^{-\frac{N-2}{N-1}}, \quad (14)$$

as the probability distribution function of self-protection in a mixed Nash equilibrium.

Case $c \leq pL$: Now let us consider a game with insurance under the more reasonable assumption $c \leq pL$; that is, insurance is not overpriced compared to expected losses. Dominance arguments indicate that a Nash strategy must be of the form $(e, s) \in \{(e, 0), e \geq 0\} \cup \{(0, 1)\}$.

Let q be the probability that a player chooses strategy $(e, s) = (0, 1)$. That is, $F(0) = q$. Because insurance is independent of protection, we can reuse Eqn. (13) with the new boundary $G(0) = 1 - q$:

$$G(e) = \left((1-q)^{N-1} - \frac{b}{pL}e\right)^{\frac{1}{N-1}} \quad (15)$$

However, since we are now including self-insurance, a second condition must hold. The payoff for strategy $(e, s) = (0, 1)$ must equal the payoff for all other strategies.

Specifically, we may compare payoffs for strategies $(e, s) = (\varepsilon, 0)$ and $(e, s) = (0, 1)$ which gives, by continuity as $\varepsilon \rightarrow 0$,

$$pL(1-q)^{N-1} = c. \quad (16)$$

Together Eqs. (15) and (16) yield:

$$F(e) = 1 - G(e) = 1 - \left(\frac{c-be}{pL}\right)^{\frac{1}{N-1}},$$

which, differentiating, gives

$$f(e) = \frac{1}{N-1} \frac{b}{pL} \left(\frac{c-be}{pL}\right)^{\frac{1}{N-1}-1}. \quad (17)$$

This allows us to compute how often strategy $(e, s) = (0, 1)$ is played:

$$q = F(0) = 1 - \left(\frac{c}{pL}\right)^{\frac{1}{N-1}}. \quad (18)$$

Result 5: In the weakest-target game with an attacker of infinite strength, a mixed Nash equilibrium strategy exists. The individual's strategy is given by Eqs. (17) and (18).

Also note that, per Eqn. (17) and continuity arguments, the upper bound for protection effort is given by $e_{\max} = c/b$, which can be less than 1 when protection costs dominate insurance costs $b > c$.

Increasing number of players N : From Eqn. (18), we can directly infer that an increase in the number of participating players decreases the probability that a full self-insurance strategy is chosen. When N grows large, q tends to zero, which means that players increasingly prefer to gamble in order to find a protection level that leaves them unharmed.

4.5 Weakest target (with mitigation)

Let us assume that there exists a Nash equilibrium where $0 < K < N$ players who satisfy $e_i = e_0 = \min(e_i, e_{-i})$, while $(N - K > 0)$ players satisfy $e_i > e_0$. We can show that such an equilibrium does not exist and that players rather congregate at the highest protection level if certain conditions are met. Due to space constraints, we will only sketch the analysis of this equilibrium. By computing the partial derivatives $\partial U_i / \partial s_i$ and $\partial U_i / \partial e_i$, and discriminating among values for e_i and s_i , we get the following results.

Result 6: In contrast to the infinite strength weakest-target game we find that a pure Nash equilibrium may exist.

- *Full protection eq.:* If $b \leq c$ we find that the full protection equilibrium $(\forall i, (e_i, s_i) = (1, 0))$ is the only possible pure Nash equilibrium.
- For $b > c$ we can show that no pure Nash equilibrium exists.
- There are no pure self-insurance equilibria.

Mixed strategy equilibrium To complement this analysis we also present the mixed strategy equilibrium. The derivation is similar to the one given by Eqs. (12–18), however, with an additional substitution step. This gives the resulting distribution,

$$F(e) = 1 - \left(\frac{c-be}{pL(1-e)}\right)^{\frac{1}{N-1}}, \quad (19)$$

so that

$$f(e) = \frac{1}{N-1} \left(\frac{(b-c)pL}{pL^2(1-e)^2}\right) \left(\frac{c-be}{pL(1-e)}\right)^{-\frac{N-2}{N-1}}.$$

Interestingly, the probability of playing $(e, s) = (0, 1)$ remains

$$q = F(0) = 1 - \left(\frac{c}{pL}\right)^{\frac{1}{N-1}} \quad (20)$$

Note that if $c < b$ there is a zero probability that $e = 1$ will be chosen by any player. The upper bound for protection effort is given by $e_{\max} = c/b$.

Result 7: In the weakest-target game with an attacker of finite strength we find that a mixed Nash equilibrium strategy exists. The relevant equations are given in Eqs. (19–20).

5. IDENTIFICATION OF SOCIAL OPTIMA

Organizations and public policy actors frequently attempt to identify policies that provide the highest utility for the largest number of people. This idea has been operationalized with the social optimum analysis. It states that a system has reached the optimum when the sum of all players' utilities is maximized. That is, the social optimum is defined by the set of strategies that maximize $\sum_i U_i$. Consider N players, and denote by $\Phi(e_1, s_1, \dots, e_N, s_N)$ the aggregate utility, $\Phi(e_1, s_1, \dots, e_N, s_N) = \sum_i U_i(e_i, s_i)$. The social optimum maximizes $\Phi(s_i, e_i)$ over all possible $(s_i, e_i) \in [0, 1]^{2N}$. Because enforcing a social optimum may at times be conflicting with the optimal strategy for a given (set of) individual(s), to enforce a social optimum in practice, we may need to assume the existence of a "social planner" who essentially decides, unopposed, the strategy each player has to implement.

5.1 Total effort game

Summing the utility given by Eqn. (2) over i , we realize that $\Phi((e_i, s_i)_{i \in \{1, \dots, N\}})$ can be expressed as a function of two variables, $E = \sum_i e_i$ and $S = \sum_i s_i$. Φ is continuous and twice differentiable in E and S , and the second derivative test tells us that the only possible extrema of Φ are reached for the boundary values of E and S , that is $(E, S) \in \{0, N\}^2$. In other words, the only possible social optima are 1) passivity (for all i , $(e_i, s_i) = (0, 0)$), 2) full protection (for all i , $(e_i, s_i) = (1, 0)$), or 3) full insurance (for all i , $(e_i, s_i) = (0, 1)$). As long as one of b or c is strictly positive, a social planner will never advise agents to invest into protection and self-insurance at the same time.

By comparing the values of Φ in all three cases, we find that if $b < pL$ and $b < c$ then all agents are required to exercise maximum protection effort $(e_i, s_i) = (1, 0)$. With $c < pL$ and $c < b$ all agents will self-insure at the maximum possible $(e_i, s_i) = (0, 1)$. A social planner will not encourage players to invest in security measures if they are too expensive ($c > pL$ and $b > pL$).

Result 8: *In the total effort security game we observe that in the Nash equilibrium there is almost always too little protection effort exerted compared to the social optimum. In fact, for a wide range of parameter settings no protection equilibria exist while the social optimum prescribes protection at a very low threshold.*

- **Protection:** Except for very unbalanced parameter settings (i.e., $pL > bN$ and $c > b + pL \frac{N-1}{N}$) agents refrained from full protection. Now full protection by all agents is a viable alternative.
- **Self-insurance:** Full self-insurance now has to compete with full protection effort under a wider range of parameters.
- **Passivity:** Agents remain passive if self-insurance is too expensive ($c > pL$). However, we find a substantial difference with respect to protection behavior. Agents would selfishly refrain from protection efforts if $pL < bN$ since they would only be guaranteed the N -th part of their investments as returns. Now the social planner can ensure that all agents protect equally so that it is beneficial to protect up until $b < pL$.

5.2 Weakest link game

In the weakest link game agents are required to protect at a common effort level to be socially efficient. We compute Φ by summing Eqn. (3) over i , and can express Φ as a function of e_i , s_i and $e_0 = \min_i(e_i)$. In particular, for all i , we obtain $\partial\Phi/\partial s_i = pL(1 - e_0) - c$. Studying the sign of $\partial\Phi/\partial s_i$ as a function of e_0 tells us that, if $b < c$ and $b < pL$ the social planner requires all

agents to protect with maximum effort $(e_i, s_i) = (1, 0)$. If $c < b$ and $c < pL$ the social planner requires all agents to self-insure $(e_i, s_i) = (0, 1)$. Finally, the Nash equilibrium and social optimum coincide when security costs are high. Agents do not invest in protection or self-insurance if $b > pL$ or $c > pL$.

Result 9: *The availability of self-insurance lowers the risk of below-optimal security in the Nash equilibrium since agents have an alternative to the unstable Pareto-optimal protection equilibrium. From the analysis of the weakest link game with many agents we know that deviation from the Pareto-optimal highest protection level is very likely. A social planner can overcome these coordination problems.*

- **Protection:** The Pareto-optimal Nash equilibrium coincides with socially optimal protection. However, the protection level would likely be lower in the Nash case due to coordination problems.
- **Self-insurance:** The self-insurance equilibria are equivalent for the Nash and social optimum analysis.
- **Passivity:** A social planner cannot expand the range of parameter values at which it would be socially beneficial to protect or self-insure while passivity would be prescribed in the Nash equilibrium.

5.3 Best shot game

We compute the social optimum by summing U_i given in Eqn. (4) over i , yielding that Φ can be expressed as a function of e_i , s_i , and e^* . It is immediate that, to maximize Φ , one should pick $e_i = 0$ for all i , except for one participant j , where $e_j = e^* \geq 0$. We then get $\partial\Phi/\partial s_i = pL(1 - e^*) - c$, which tells us under which conditions on e^* (and consequently on b , c , and pL) self-insurance is desirable.

We find that if $b/c < N$ (i.e., protection is not at a prohibitive cost compared to insurance and/or there is a reasonably large number of players), the social optimum is to have one player protect as much as possible, the others not protect at all, and no one insures. In practice, this may describe a situation where all participants are safely protected behind an extremely secure firewall. If, on the other hand $b/c > N$, which means there are either few players, insurance is very cheap compared to protection, then the best strategy is to simply insure all players as much as possible.

Result 10: *In the best shot security Nash outcome there is almost always too little effort exerted compared to the social optimum. Exceptions are few points in which full self-insurance remains desirable for the social planner and all agents remain passive.*

- **Protection:** Surprisingly, while protection is not even a Nash strategy we find that a social planner would elect an individual to exercise full protection effort.
- **Self-insurance:** Full self-insurance by every player is only desirable if protection costs are large. Therefore, for most cases the strategy of a social planner will not coincide with the only Nash equilibrium strategy.
- **Passivity:** In the Nash equilibrium agents are also too inactive. Passivity is highly undesirable from a social planner's perspective. Only if $NpL < b$ no agent will be selected to exercise maximum protection effort (while self-insurance might remain an option).

It is important to note that the social optimum variation that requires full protection by one individual results in the whole population being unharmed, since one highly secure individual is enough to thwart all attacks. Therefore, it is easy to see that protection is extremely desirable from a planners perspective. Out of the three classical public goods games with homogeneous agents the best shot game can benefit the most from a guiding hand.

5.4 Weakest-target security game (without mitigation)

We compute the social optimum by using Eqn. (8), assuming that $1 \leq K \leq N$ players pick $e_0 = \min_i(e_i)$. By studying the variations on Φ as a function as e_i , as a function of K , and as a function of s_i (for both the K players picking e_0 and the remainder of the players), we find that in the weakest-target game without mitigation a social planner would direct a single player to exacerbate no protection effort.

Essentially, this player serves as a direct target for a potential attacker. However, as long as $c < pL$ the player would be directed to maximize self-insurance $(e_i, s_i) = (0, 1)$. If insurance is too expensive ($c > pL$) then the social planner would prefer to leave the player uninsured $(e_i, s_i) = (0, 0)$. This strategy is independent of the cost of protection. The remaining $N - 1$ players have to select their protection effort as $e_i = \varepsilon > 0$ (as small as possible). These players will not be attacked, and therefore will set their self-insurance to the possible minimum $(\varepsilon, 0)$. Passivity by all players is never an option in the social optimum.

Result 11: *A social planner can easily devise a strategy to overcome the coordination problems observed in the Nash analysis for the weakest-target game with mitigation. We found that no pure Nash strategy exists and, therefore, had to rely on the increased rationality requirement for entities to play a mixed strategy.⁴ The average payoff for each player in the social optimum is considerably higher compared to the mixed Nash equilibrium.*

Understandably, without side-payments the node with the lowest protection effort is worse off compared to his peers. However, the social planner could choose to devise a so-called “honeypot” system with the sole goal of attracting the attacker while only suffering a marginal loss. A honeypot is a computer system (or another device) that is explicitly designed to attract and to be compromised by attackers. It serves usually a double purpose. First, it will distract attention from more valuable targets on the same network. Second, if carefully monitored it allows gathering of information about attacker strategies and behaviors, e.g., early warnings about new attack and exploitation trends [30].

An interesting aspect of the social optimum solution is the question how the individual is selected (if a honeypot system cannot be devised). Obviously, a social planner might be able to direct an individual to serve as a target (in particular, if $c < pL$). However, if insurance costs are large being a target requires an almost certain sacrifice (dependent on the value of p). In anthropology and economics there are several theories that relate to an individuals willingness to serve as a sacrificial lamb. Most prominently, altruism and heroism come to mind. Simon also introduced the concept of docility. This theory refers to an individual’s willingness to be targeted or to defer to the superior knowledge of others [39].

⁴Economists are generally cautious regarding the assumption that individuals can detect and adequately respond to mixed strategy play by opponents [36].

5.5 Weakest-target security game (with mitigation)

We adopt the same strategy for finding Φ ’s maximum as in the unmitigated case – that is, summing Eqn. (6) over i , and then studying the variations of Φ over K , s_i and e_0 .

The first observation is that the social planner might prescribe the same strategy as in the case of the weakest-target game without mitigation. However, now the planner has a second alternative. Since an attacker will not be able to compromise players if they are fully protected we find that $(e_i, s_i) = (1, 0)$ for all N players is a feasible strategy. The tipping point between the two strategies is at $Nb < c$. If this condition holds the social planner would elect to protect all machines in favor of offering one node as honeypot and investing in its self-insurance. Note that again we find that if protection and self-insurance are extremely costly the planner will elect to sacrifice one entity without insurance. Passivity is not a preferable option.

Result 12: *Compared to the weakest-target game without mitigation the social planner is better off if protection is cheap. Otherwise the planner has to sacrifice a node with or without self-insurance. Interestingly, while compared to the pure Nash equilibrium outcome the social planner can increase the overall utility in the network we find that security expenditures are lowered. In the Nash equilibrium agents were willing to fully protect against threats as long as $(b \leq c)$.*

*The last observation also holds for the mixed strategy case in both weakest-target games (with or without mitigation). That is, agents exert **more** effort in the Nash equilibrium (except when $Nb < c$ for the game with mitigation).*

6. DISCUSSION OF RESULTS

The results we obtained, and notably the disconnect between social optima and Nash equilibria we observed, lead to a number of remarks that may prove relevant to organizational strategy. However, we want to preface this discussion by pointing out that our analysis is a first comparison of different security games with two security options under common, but restrictive assumptions.

Most notably, we assume agents to be risk-neutral providers of the public protection good. In our game formulation we also simplified cost of protection (and insurance) to be linear. Including different risk preferences, as well as uncertainty and limited information about important parameters of the game would be important steps towards a sensitivity analysis of our results. Shogren found, for example, that risk-averse agents will increase their contributions if information about other agents actions is suppressed [38]. Others, e.g., [34], have obtained more nuanced results. We defer a more extensive analysis of such phenomena to future work, but believe that the main trends and differentiating features between security games we observed remain largely unchanged.

Security scenario identification: We find that security predictions vary widely between the five different games. Similarly, policies set by a social planner do not only yield different contribution levels but may also switch the recommended security action from protection to self-insurance and vice versa. Chief Security Officers’ tasks involve a careful assessment of threat models the company is faced with.

We want to emphasize that an integral part of the threat model should be an assessment of the organizational structure including system resources and employees. Similarly important is a detailed consideration whether resources are protected independently or by an overarching system policy. For example, replication, redundancy and failover systems (that automatically switch to a standby

database, server or network if the primary system fails or is temporarily shut down for servicing) should most likely not be treated as independent resources.

Managers should consider how the organizational structure of resources matches potentially existing policies. For example, we can see that a policy that requires full protection by every individual is sub-optimal if the most likely threat and organizational structure fits the description of a best shot game. Contributions resources are squandered and are likely to deteriorate. Not to mention that employees may simply ignore the policy over time. See, for example, recent survey results that highlight that 35% of white-collar employees admit to violations of security policies [25].

Security scenario selection: A security professional might be faced with a unidentifiable organization and system-policy structure. However, we want to highlight that our research allows a more careful choice between security options if managers can redesign organizations and policies. For example, the choice between a system-wide firewall and intrusion detection system versus an individual alternative has important implications on how incentives drive security-relevant behavior over time. Individual systems will better preserve incentives, however, might have negative cost implications. The same choice applies between the availability of backup tools and protective measures.

Leveraging strategic uncertainty: The example of the weakest-target game shows the importance of the degree of dependency between agents. We show that in larger organizations a much lower average level of self-insurance investments will be achieved because the strategic dependence between actors is reduced. However, in turn more agents will elect to protect their resources ($e_i > 0$ for more players). In contrast, agents in small groups will respond to the increasing strategic uncertainty caused by the increased interdependency by self-insuring their resources more often.

Introducing a social planner into the weakest-target game completely removes strategic uncertainty and leads to both reduced self-insurance and protection investments. This apparent paradox emphasizes that higher security investments do not necessarily translate in higher security – but instead that *how* the investments are made are crucial to the returns.

7. CONCLUSIONS

We consider the problem of decision-making with respect to information security investments. To that effect, we model security interactions through a careful selection of games, some established (weakest-link, best-shot, and total effort) and some novel (weakest-target, with or without mitigation). All of these games offer players two independent decision parameters: a protection level, e , which determines the level of security a player chooses for his resources; and a self-insurance level, s , which mitigates losses, shall a successful attack occur. We postulate that the five games considered cover a vast majority of practical security situations, and study them both from a rational agent's perspective (Nash equilibrium analysis) and from a central planner's view (social optimum analysis).

Our main findings are that the effects of central planning compared to laissez-faire considerably differ according to the game considered. While in a number of traditional cases borrowed from the public good literature, we observe that a central planner may increase the average protection level of the network, we also note that strategic decisions are highly impacted by the level of interdependency between the actions of different players.

In particular, we found that the common wisdom that having a central planner who decides upon security implementation always yields higher protection contributions by individual players does

not hold. Indeed, it may at times be much more advantageous from an economic standpoint to invest in self-insurance instead of protecting systems, or to select a few, unprotected, sacrificial lambs in order to divert the attention of potential attackers. This is particularly the case in situations which exhibit a “strategic uncertainty” due to a very strong correlation between the actions of different agents, for instance, in our weakest-target game where the least secure player is always the one attacked.

7.1 Future research directions

The work presented here opens a number of avenues for future research. First, we have looked at homogeneous populations of users, where all participants have the same utility function. In practice, this homogeneity assumption is reasonable in a number of important cases, particularly when dealing with very large systems where a large majority of the population have the same aspirations. For instance, most Internet home users are expected to have vastly similar expectations and identical technological resources at their disposal; likewise, modern distributed systems, e.g., peer-to-peer or sensor networks generally treat their larger user base as equals. It will be nevertheless prudent to study whether considering heterogeneous user populations impacts the results we obtained here and in what way. Varian [43], for instance, evidences important asymmetric user behaviors due to heterogeneity in his reliability games. We also plan to extend our work to explore the impact of fines and liability rules on security investments [9, 16].

Second, this paper assumes that execution of a player's strategy is always perfect, and that all players are perfectly rational. As has been discussed elsewhere, e.g., [12], this assumption generally leads to idealized models, which deserve to be complemented by empirical studies. In that respect, we are currently developing a set of laboratory experiments to conduct user studies and attempt to measure the differences between perfectly rational behavior and actual strategies played. Our preliminary investigations in the field notably evidence that players often experiment with different strategies to try to gain a better understanding of the game they are playing.

Reconciling observed experimental behavior and theoretical analysis with the design of meaningful security policies is a very challenging goal. We do hope that the present paper will encourage research in this area.

8. ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable comments and editorial guidance. Paul Laskowski greatly improved this manuscript with his tremendously helpful feedback. This work is supported in part by the National Science Foundation under ITR award ANI-0331659. Jens Grossklags' research is also partially funded by TRUST (Team for Research in Ubiquitous Secure Technology), under support from the NSF (award CCF-0424422) and the following organizations: BT, Cisco, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, Telecom Italia, United Technologies, and AFOSR (#FA 9550-06-1-0244).

9. REFERENCES

- [1] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, January–February 2005.
- [2] E. Adar and B. Huberman. Free riding on Gnutella. *First Monday*, 5(10), Oct. 2000.
- [3] R. Anderson. Why cryptosystems fail. In *Proc. ACM CCS'93*, pages 215–227, Fairfax, VA, Nov. 1993.

- [4] R. Anderson. Why information security is hard – an economic perspective. In *Proc. ACSAC'01*, New Orleans, LA, Dec. 2001.
- [5] R. Anderson and T. Moore. The economics of information security. *Science*, 314(5799):610–613, Oct. 1998.
- [6] AOL/NSCA. Online safety study, Dec. 2005. http://www.staysafeonline.org/pdf/safety_study_2005.pdf.
- [7] T. August and T. Tunca. Network software security and user incentives. *Mgmt. Science*, 52(11):1703–1720, Nov. 2006.
- [8] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. In *Proc. (online) WEIS'06*, Cambridge, UK, June 2006.
- [9] J. Brown. Toward an economic theory of liability. *Journal of Legal Studies*, 2(2):323–349, June 1973.
- [10] Bruskin Research. Nearly one in four computer users have lost content to blackouts, viruses and hackers according to new national survey, 2001. http://www.corporate-ir.net/ireye/ir_site.zhtml?ticker=iom&script=410&layout=-6&item_id=163653.
- [11] J. Bull, L. Gong, and K. Sollins. Towards security in an open systems federation. In *Proc. ESORICS'92, Springer LNCS No. 648*, pages 3–20, Toulouse, France, Nov. 1992.
- [12] N. Christin, J. Grossklags, and J. Chuang. Near rationality and competitive equilibria in networked systems. In *Proc. ACM SIGCOMM'04 PINS Workshop*, pages 213–219, Portland, OR, Aug. 2004.
- [13] D. Clark, J. Wroclawski, K. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow's Internet. In *Proc. ACM SIGCOMM'02*, pages 347–356, Pittsburgh, PA, Aug. 2002.
- [14] G. Danezis and R. Anderson. The economics of resisting censorship. *IEEE Security & Privacy*, 3(1):45–50, January–February 2005.
- [15] I. Ehrlich and G. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, July 1972.
- [16] E. Fehr and S. Gaechter. Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994, Sept. 2000.
- [17] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of Internet miscreants. In *Proc. ACM CCS'07*, Alexandria, VA, Oct./Nov. 2007.
- [18] L. Gordon and M. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, Nov. 2002.
- [19] S. Gordon. The generic virus writer. In *Proc. Intl. Virus Bulletin Conf.*, pages 121 – 138, Jersey, Channel Islands, 1994.
- [20] S. Gordon. Virus writers - the end of the innocence? In *10th Annual Virus Bulletin Conference (VB2000)*, Orlando, FL, Sept. 2000. <http://www.research.ibm.com/antivirus/SciPapers/VB2000SG.htm>.
- [21] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, Dec. 1968.
- [22] K. Hausken. Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. *Information Systems Frontiers*, 8(5):338–349, Dec. 2006.
- [23] J. Hirshleifer. From weakest-link to best-shot: the voluntary provision of public goods. *Public Choice*, 41(3):371–386, Jan. 1983.
- [24] P. Honeyman, G. Schwartz, and A. van Assche. Interdependence of reliability and security. In *Proc. (online) WEIS'07*, Pittsburgh, PA, June 2007.
- [25] Information Systems Audit and Control Association. Telephone survey conducted by MARC Research, Oct. 2007. <http://biz.yahoo.com/bw/071031/20071031005079.html?v=1>.
- [26] H. Kunreuther and G. Heal. Interdependent security. *J. Risk and Uncertainty*, 26(2–3):231–249, Mar. 2003.
- [27] S. Malphrus. The “I Love You” computer virus and the financial services industry, May 2000. <http://www.federalreserve.gov/BoardDocs/testimony/2000/20000518.htm>.
- [28] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer worm. *IEEE Security and Privacy*, 1(4):33–39, July 2003.
- [29] D. Moore, C. Shannon, and J. Brown. Code-Red: a case study on the spread and victims of an internet worm. In *Proc. ACM/USENIX IMW'02*, pages 273–284, Marseille, France, Nov. 2002.
- [30] N. Provos. A virtual honeypot framework. In *Proc. USENIX Security'04*, pages 1–14, San Diego, CA, Aug. 2004.
- [31] E. Rescorla. Security holes... who cares? In *Proc. USENIX Security'03*, pages 75–90, Washington, DC, Aug. 2003.
- [32] J. Saltzer, D. Reed, and D. Clark. End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4):277–288, Nov. 1984.
- [33] T. Sandler and K. Hartley. Economics of alliances: The lessons for collective action. *Journal of Economic Literature*, XXXIX(3):869–896, Sept. 2001.
- [34] T. Sandler, F. Sterbenz, and J. Posnett. Free riding and uncertainty. *Economic Review*, 31(8):1605–1617, Dec. 1987.
- [35] T. Schelling. *The Strategy of Conflict*. Oxford University Press, Oxford, UK, 1965.
- [36] J. Shachat and J. Swarthout. Do we detect and exploit mixed strategy play by opponents? *Mathematical Methods of Operations Research*, 59(3):359–373, July 2004.
- [37] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Trans. Networking*, 3(6):819–831, Dec. 1995.
- [38] J. Shogren. On increased risk and the voluntary provision of public goods. *Social Choice and Welfare*, 7(3):221–229, Sept. 1990.
- [39] H. Simon. Altruism and economics. *American Economic Review*, 83(2):156–161, May 1993.
- [40] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In *Proc. ACM EC'01*, pages 38–47, Tampa, FL, Oct. 2001.
- [41] The HoneyNet Project. Know your enemy: the tools and methodologies of the script-kiddie, July 2000. <http://project.honeynet.org/papers/enemy/>.
- [42] J. Van Huyck, R. Battalio, and R. Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–248, 1990.
- [43] H. Varian. System reliability and free riding. In L. Camp and S. Lewis (ed.), *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.