Security and Insurance Management in Networks with Heterogeneous Agents

Jens Grossklags
School of Information
UC Berkeley
jensg@sims.berkeley.edu

Nicolas Christin
INI/CyLab Japan
Carnegie Mellon University
nicolasc@cmu.edu

John Chuang
School of Information
UC Berkeley
chuang@sims.berkeley.edu

ABSTRACT

Computer users express a strong desire to prevent attacks and to reduce the losses from computer and information security breaches. However, security compromises are common and widespread and highly damaging. Next to attackers' increased sophistication, a root cause for the harm inflicted is that users often fail to optimally protect their resources or to recover gracefully from a security breach.

We argue that users often underestimate the strong mutual dependence between their security strategies and the economic environment (e.g., threat model) in which these choices are made and evaluated. This misunderstanding weakens the effectiveness of users' security investments, and is compounded by heterogeneity within the user population, in some cases further reducing incentives for cooperation and coordination.

We study how economic agents invest into security in five different economic environments, that are characteristic of different threat models. We consider generalized models of traditional public goods games (e.g., total effort and weakest link) and two recently proposed games (e.g., weakest target game). Agents may split their contributions between a public good (protection) and a private good (self-insurance).

Our analysis centers on how agents respond to incentives when important parameters of the game (i.e., loss probability, loss magnitude, and cost of technology) are heterogeneous in the agent population. We also highlight key differences to the case of homogeneous decision makers. For example, security investments may become substantially more sensitive to the size of the network. We extend our results to discuss important modes of intervention.

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, Security, Reliability

EC'08, July 8–12, 2008, Chicago, Illinois, USA.
Copyright 2008 ACM 978-1-60558-169-9/08/07 ...$5.00.

Keywords

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

1. INTRODUCTION

A majority of computer and network users indicate interest in preventing attacks and limiting the damages from security breaches [1]. At the same time, measurement studies and surveys [4, 7, 36] show strong evidence that security precautions, be they patching, spyware-removal tools, or even sound backup strategies, are absent from a vast majority of systems surveyed.

As a result, security problems seem to be multiplying – see for instance [28, 29, 30] for some of the cases most discussed in popular media. Threats include attacks on the network as a whole, attacks on selected end-points, racketeering and/or blackmail linked to distributed denial of service, undesirable forms of interactions such as spam e-mail, and annoyances such as Web pages that are unavailable or defaced.

While precise estimates of the impact of each of these threats are still subject to debate, a somewhat general consensus is that taken all together, security incidents have caused billions of dollars in damages, and could easily be dwarfed by future losses, if the situation is not addressed quickly [41]. Damages include loss of information, incapacity to conduct business for extended periods of time, productivity losses, and surcharge in repair and maintenance, to name a few. The situation has degraded to the point that underground markets capitalizing on security issues are emerging [15, 24].

It therefore becomes urgent to understand why most individuals and corporations either do not implement sufficient security on their systems, or invest in the wrong things, when, by and large, security technology is readily available, and most end-users have a stated interest in improving the security of the systems they own or operate.

To obtain elements of answers to the above research question, we focus on an analysis of five stylized "security games," with a small number of decision parameters upon which each user can act. These five games provide a series of simple abstractions that can capture the essence of most of the security interactions observed in the field [19].

To that effect, we build upon public goods literature [22, 39], and consider the classical best shot, total effort, and weakest link games in a security context with heterogeneous agents. We also revisit our proposed "weakest target" game [19], which allows us to describe a whole class of attacks ranging from insider threats to very aggressive worms. Compared to prior work, the key novelty of the model behind our proposed security games is to decouple protection investments (e.g., setting up a firewall) from insurance coverage (e.g., archiving data as back up). In short, our work is the first to tackle security as a hybrid between a private and a public good.

The first contribution of the present paper is to extend and generalize models we previously proposed for homogeneous agents [19] to the significantly more complex heterogeneous agent case.

The second contribution of this paper is to exploit the results from the analysis to evaluate the impact of possible (centralized or distributed) intervention policies aiming at reaching an outcome beneficial to society as a whole.

We point out that this study relies on game theory, mostly using Nash equilibrium concepts, which may present some limitations [11]. We nevertheless postulate that the models and results derived here allow to gain valuable insights into user behavior and policy impact, which we ultimately plan on validating through behavioral user studies.

The rest of this paper is organized as follows. We elaborate in Section 2 on the relationship of our work with related research, and extend our game-theoretic models to take into account agent heterogeneity in Section 3. We analyze Nash equilibria stemming from these games in Section 4, and use this analysis to look into possible intervention mechanisms in Section 5. We conclude in Section 6.

2. BACKGROUND AND MOTIVATION

Using economic tools to aid in security analysis has gained increasing relevance over the past few years. Indeed, attackers have become more and more akin to rational economic actors, e.g., motivated by greed [15]. In addition, the democratization of information networks has led to numerous business opportunities, which have in turn translated in a change in the behavior of network participants. They are indeed becoming increasingly self-interested, and increasingly view information networks as competitive markets rather than cooperative platforms [12].

2.1 Economics of security

The key findings coming from studying network security through the prism of economics are 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 [2, 3, 26].

A large body of research [5, 18, 21, 34] has been devoted to analyzing optimal security investments from an individual choice's perspective. For instance, some of these works attempt to characterize optimal patching strategies at a end-host when patching itself is a costly operation. As such, this research considers agents as acting in isolation, in response to a given exogenous threat. Our work, on the other hand, considers strategic interactions when players face network effects, that is, when each individual choice affects the security of the whole network. While players are not necessarily hostile to each other (although they may be), they usually do not have the best interest of the whole network in mind when making their security choices.

2.2 Security as hybrid goods

Our work builds on the public goods literature [22, 39] by treating security as a hybrid between public and private goods. We identify two key components of a security strategy: self-protection (e.g., patching system vulnerabilities) and self-insurance (e.g., having good backups) [19].

Self-protection denotes the ability to reduce the probability of a loss. Self-insurance, on the other hand, characterizes a reduction in the magnitude of a loss, when a disastrous event occurs. Similar to the exposition by Varian [39], individual self-protection strategies (chosen by each player) have an effect on the protection levels of all players in the network, which is a key characteristic of public goods. Self-insurance only affects the player who subscribes to it, and is consequently a private good.

Our analysis complements the work of Ehrlich and Becker, who were the first to consider the joint concepts of self-protection and self-insurance [14], by extending the discussion to the public goods and security context.

2.3 Heterogeneity in system security

In a stark contrast to our previous work [19], the salient feature of the research presented in this paper is to consider security as a combination of private and public goods in the context of *heterogeneous* agents.

Both the homogeneous and heterogeneous cases are relevant to security analysis. Homogeneous agents are characteristic of large populations following the same practices and choices by end-users, for instance, when most security decisions (e.g., patching) are automated, and all users run similar software. The lack of diversity, in particular in the market for operating systems, lends credibility to such scenarios [17], and is cited as a strong motivator for developers of malicious code to exploit the resulting correlated risks or to cheaply repeat attacks.

However, there are strong reasons to compare our earlier findings with a model that includes heterogeneous agents into a model of security decision making.

Security through diversity. Recent technical proposals aim to achieve higher resilience to attacks by introducing diversity in network and protocol design. For example, Zhuang et al. report of a set of formal analysis tools that introduce heterogeneity in multi-person communication protocols [42]. O'Donnell and Sethu develop and test distributed algorithms optimizing the distribution of distinct software modules to different nodes in a network [31]. Research in IT economics has evaluated the decision making of a firm when faced with the option of increased diversity in its software base. In Chen et al., the decision for increased heterogeneity depends largely on the assumed risk attitudes of the organization [10]. Investments into heterogeneity will change the expectation of losses and attack probabilities, but they also impact the cost of protection and self-insurance.

Chameleonic threats. Increased diversity is not a sufficiently strong protection against correlated security threats anymore. Already in 1995 the first macro viruses started targeting MS office on all compatible systems.[1] Modern cross-platform malware is capable of targeting also different operating systems. For instance, Linux-Bi-A/Win-Bi-A is written in assembler and able to compromise Windows and Linux platforms. Malicious code is also capable of crossing the boundary between desktop and mobile devices. For example, the hybrid pathogen Nimda, a worm that can spread as a virus as well, has successfully propagated on different media such as floppies, portable hard drives, and USB pen drives [40].

Potentially even more disruptive is malware carrying multiple exploit codes at once. For example, Provos et al. report that Web-based malware often includes exploits that are used 'in tandem' to download, store and then execute a malware binary [32]. These trends render users vulnerable to propagated threats if owners of different IT systems perceive protection as too costly or ineffective.

[1] The macro virus (Winword-Concept) targeted Microsoft Word on Apple and Microsoft systems. For more details see: `http://web.textfiles.com/virus/macro003.txt`.

Heterogeneous investments patterns. Different organizations follow distinct patterns of IT investment. Parts of organizations often depend on legacy systems including weakly protected systems, or "boat anchors" with limited value to an organization [41]. Such legacy systems can allow skilled attackers to intrude a network. More generally, organizations and end users justify security investments with different assumptions about potential losses and probabilities of being attacked. This often depends on different knowledge about threats and means of protection and insurance [1]. This diversity is reflected in users' choices and security practices [4, 7]. Similarly, security decisions can follow different security paradigms often reflected in different organizational structures, for instance remote replication vs. offsite tape storage.

Finally, heterogeneous agents have notable implications in terms of policy design. For instance, Bull et al. [8] observe the state of heterogeneous networks and argue that no single security policy will be applicable to all circumstances. They argue 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.

In this paper, we formally explore such theses, by studying individuals' incentives in non-cooperative games. In particular, we focus on the impact of heterogeneous agents on system security in different network structures.

3. FIVE CANONICAL SECURITY GAMES

A security game is a game-theoretic model that captures essential characteristics of decision making to protect and self-insure resources within a network of agents. In this section, we summarize the security games we analyze, and extend models we previously proposed [19] to the heterogeneous agent case.

As discussed earlier, we model security as a hybrid between public and private goods. On the one hand, as was previously observed by Varian [39], the success of security (or reliability) decision making frequently depends on a joint protection level determined by all participants of a network. The computation of the protection level will often take the form of a public goods contribution function. Because network protection is a public good, it may allow, for certain types of contribution functions, individuals to free-ride on others' efforts. At the same time, some individuals may also suffer from inadequate protection efforts by other members if those have a decisive impact on the overall protection level.

In addition to self-protection, network participants 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.

All security games we introduce share the following key assumptions: (i) all entities in the network share a single purely public protection output, and (ii) a single individual decides on protection efforts for each entity – we do not assume a second layer of organizational decision making.

Different from our previous exposition [19], protection costs per unit are not necessarily identical for each entity, and, while in the formal analysis that follows we make the assumption that all decisions are made simultaneously, we later discuss the impact of relaxing the synchronization assumption.

We develop security games from a basic model with the following payoff structure. Each of $N \in \mathbb{N}$ players receives an individual endowment M_i. If she is attacked and compromised successfully she faces a loss L_i. Attacks arrive with a probability of p_i ($0 \leq p_i \leq 1$), which albeit exogenous, is also dependent on the player under consideration; p_i remains constant over time.[2] 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, $b_i \geq 0$ and $c_i \geq 0$ denote the unit cost of protection and insurance, respectively. The generic utility function of Player i is defined as:

$$U_i = M_i - p_i L_i (1 - s_i)(1 - H(e_i, e_{-i})) - b_i e_i - c_i s_i \ , \quad (1)$$

where, following common game-theoretic notation, e_{-i} denotes the set of protection levels chosen by players other than i. H is a contribution function of e_i, which is required to 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_i), 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 rely on 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., [22, 35, 39]). 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 games, which we originally introduced only in the context of homogeneous agents [19].

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

$$U_i = M_i - p_i L_i (1 - s_i)(1 - \frac{1}{N} \sum_k e_k) - b_i e_i - c_i s_i \ . \quad (2)$$

Economists identified the sum of efforts (or total effort) contribution function long before the remaining cases included in this paper [22]. We consider a slight variation of this game to normalize it to the desired parameter range.

As a practical example of a total effort game in practice, 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, which 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:

$$U_i = M_i - p_i L_i (1 - s_i)(1 - \min(e_i, e_{-i})) - b_i e_i - c_i s_i \ . \quad (3)$$

The weakest-link game is the most often recognized public goods problem in computer security. Once the perimeter of an organization is breached it is often possible for attackers to leverage this

[2] As will become clear through our mathematical exposition, rather than p_i or L_i, the quantity that drives the computation of the various equilibria is the expected loss due to attacks, $p_i L_i$. Hence, the results we derive here will be identical to those we would obtain if we had $p = p_i$ for all i.

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. Another example is that of a two-way communication (e.g., TCP flow), where the security of the communication is determined by the least secure of the communication parties. For instance, a TCP flow between a host with a perfectly secure TCP/IP stack and a host with an insecure TCP/IP stack can be easily compromised.

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

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

Among information systems, networks with built-in redundancy share resilience qualities with the best shot security game; for instance, to completely take down communications between two (presumably highly connected and highly secure) backbone nodes on the Internet, one has to shut down all possible routes between these two nodes. Other examples of such networks with built-in redundancy include peer-to-peer networks or sensor networks. Censorship-resistant networks are another instance of best shot games. A piece of information will remain available to the public domain as long as a single node serving that piece of information can remain unharmed [13].

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 [11]. 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

$$U_i = \begin{cases} M_i - p_i L_i (1 - s_i) - b_i e_i - c_i s_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M_i - b_i e_i - c_i s_i & \text{otherwise.} \end{cases} \quad (6)$$

The weakest target game 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 [38]; or, when the attacker's goal is to commandeer the largest number of machines using the smallest investment possible [15]. 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

$$U_i = \begin{cases} M - p_i L_i (1 - s_i)(1 - e_i) - b_i e_i - c_i s_i & \text{if } e_i = \min(e_i, e_{-i}), \\ M - b_i e_i - c_i s_i & \text{otherwise.} \end{cases} \quad (8)$$

This game represents a nuanced version of the weakest target game. Here, an 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: Generalizing the above games, as we do in this paper, to heterogeneous players, still shares some limitations of the homogeneous case [19]. For instance, as hinted by Hirshleifer [22], practical scenarios may involve social composition functions combining two or more of these five games. Revisiting our earlier examples, protecting a communication flow between two hosts may be a "weakest-link" game, until a certain level of host security is reached at both hosts. At that point, the attacker may start to target the routes between the hosts rather than the hosts themselves, and it becomes a "best-shot" game. Other realistic environments may be better characterized by slight variations on a given game (e.g., "the total of the three best shots"). We nevertheless believe that the five games described above may capture a large number of practical cases, as our argument made in earlier work [19] is actually strengthened by extending the models to heterogeneous players.

4. NASH EQUILIBRIUM ANALYSIS

In this section, we derive Nash equilibria for the five different cases of security games. Our focus is to understand how the inclusion of heterogeneous actors influences predictions compared to a model with representative agents [19]. In Section 2, we have discussed arguments for and against homogeneity in security models. In the modeling of economic phenomena, added complexity (e.g., adding agents with more diverse tastes) does not always change strategic predictions substantially. On the other hand, we expect that heterogeneity impacts the actions of agents in security games in different ways, for example by: 1) Negotiating the trade-off between protection and insurance, 2) Highlighting certain strategies and focal points due to the inherent differences in the agent population, 3) (De-)stabilizing equilibrium predictions derived in the homogeneous case. We expect several conclusions from the homogeneous case to remain relevant. But as Hartley [20] argued "representative agents models conceal heterogeneity whether it is important or not." This analysis aims at pinpointing key differences and discuss their implications.

4.1 Total effort

The total effort game yields considerably different results depending on the number of players involved.

Two-player game Let us first start the discussion for the simple case $N = 2$. From the game description given by Eqn. (2), we get $U_1(e_1, s_1) = M_1 - p_1 L_1 (1 - s_1)(1 - (e_1 + e_2)/2) - b_1 e_1 - c_1 s_1$ for Player 1. The second partial derivative test indicates that there is no local extremum, so that the only possible maxima of U_1 are given by $U_1(0, 0) = M_1 - p_1 L_1 (1 - e_2/2)$, $U_1(1, 0) = M_1 - p_1 L_1 (1/2 - e_2/2) - b_1$, $U_1(0, 1) = M_1 - c_1$, or $U_1(1, 1) = M_1 - b_1 - c_1$. With $b_1 > 0$, we immediately see that $U_1(0, 1) >$

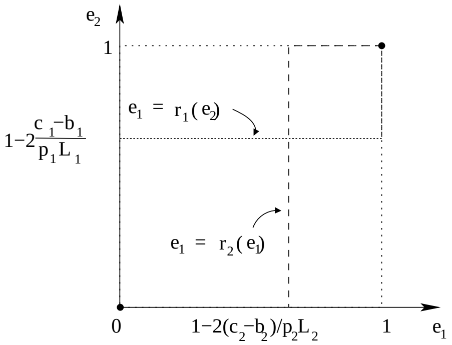

Figure 1: Reaction functions for a two-player total effort game. Bold lines and dots indicate potential Nash equilibria.

$U_1(1,1)$, which tells us that fully insuring and protecting at the same time is a strictly dominated strategy for Player 1. The passivity strategy $(e_i, s_i) = (0,0)$ dominates the "protect-only" strategy $(e_i, s_i) = (1,0)$ when $b_1 > p_1L_1/2$.

Assuming $b_1 \le p_1L_1/2$, the "protect-only" $(1,0)$ strategy dominates the "insure-only" $(0,1)$ strategy for Player 1 if and only if (all quantities being assumed to be defined):

$$e_2 > 1 - 2\frac{c_1 - b_1}{p_1L_1} \ . \tag{9}$$

A similar rationale yields the corresponding conditions for Player 2, leading to the reaction functions $e_1 = r_1(e_2)$ and $e_2 = r_2(e_1)$ plotted in Figure 1. By definition, Nash equilibria are characterized by fixed points $e_1 = r_1(e_2) = e_2 = r_2(e_1)$. From the above analysis summarized in Figure 1, this occurs for two values: when both agents fully protect and when both agents abstain from investing in protection. We note that both fixed points are stable, meaning that, if they are reached, minimal deviations in the strategy of one player are unlikely to perturb the actions of the other player.

Result 1: *The two-player total effort security game with heterogeneous agents presents the following equilibria:*

- *Full protection eq.:* If $b_1 \le p_1L_1/2$, $b_2 \le p_2L_2/2$ (protection costs are modest for both players), and the initial values $e_1(0)$ and $e_2(0)$ satisfy either $e_1(0) > 1-2(c_2-b_2)/(p_2L_2)$ or $e_2(0) > 1 - 2(c_1 - b_1)/(p_1L_1)$ (at least one player is initially fairly secure, or at least one player faces very high insurance costs) then the (only) Nash equilibrium is defined by both players protecting but not insuring, that is, $(e_i, s_i) = (1,0)$.

- *Multiple eq. without protection:* If the conditions above do not hold, then we have an insecure equilibria. Both players converge to $e_1 = 0$ and $e_2 = 0$. Their respective investments in insurance depend on whether their insurance premium is smaller than their potential losses: a player will fully insure if and only if $c_i < p_iL_i$, and will be passive otherwise.

A particularly interesting feature of the two-player version of the game is that expensive insurance or protection costs at *either* of the players directly condition which equilibrium can be reached. For instance, if one of the players has to pay a very high insurance premium in front of its protection costs, she will elect to protect, likely leading the other player to protect as well. Conversely, if either of the players faces a high protection premium ($b_i > p_iL_i/2$), the game will likely converge to an equilibrium without protection efforts. As we discuss later, this property can be used by some form of intervention to have the game converge to a desirable equilibrium.

More generally, in this game, each of the two players generally tracks what the other is doing. When moves are made perfectly simultaneously, this may result in oscillations between insecure and secure configurations. The only exception to this tracking behavior occurs when one player faces high security costs and a low insurance premium, while the other faces the opposite situation (low security costs, very high insurance premium). In such a case, the game converges to the first player insuring, and the second player protecting. In short, extreme parameter values allow to remove network effects in this game.

N-player game (N large) In the more general case $N \ge 2$, we first notice that, for a security strategy to be meaningful, we need to have $b_i < p_iL_i/N$. This means that, as the number of players increases, individual protection costs have to become very small, or expected losses have to considerably increase. Failing that, insurance or passivity is always a better option.

Second, from Eqn. (2), we obtain that Eqn. (9) is generalized to

$$\frac{1}{N-1}\sum_{j \ne i} e_j > 1 - \frac{N}{N-1}\frac{c_i - b_i}{p_iL_i} \ , \tag{10}$$

as a condition for player i to select a protection-only strategy as opposed to an insurance-only strategy. Eqn. (10) tells us that, for large values of N, changes in a single player's protection strategy are unlikely to have much of an effect on the other players' strategies. Indeed, each player reacts to changes in the average protection level over the $(N-1)$ other players.

This observation brings the question of exactly how robust the N-player game is to a change in the strategy played by a given individual. Are "domino effects" possible, where changes in a single player's strategy, albeit with a minimal effect on all other players, lead another player to switch strategies, and eventually to large groups changing their plays?

To help us answer this question, let us consider $N > 2$, and $K \le N$ arbitrary players that are initially (at time 0) unprotected. For instance, assume without loss of generality that Players $1, \ldots, K$ are initially unprotected, and that

$$\frac{c_2 - b_2}{p_2L_2} \ge \frac{c_3 - b_3}{p_3L_3} \ge \ldots \ge \frac{c_K - b_K}{p_KL_K} \ .$$

Further assume that at a later time $t > 0$, Player 1 switches her strategy to full protection, that is, $e_1(t) = 1$. Assuming all players may have an incentive to protect (i.e., for all i, $b_i < p_iL_i/N$), Player 2 would also switch to full protection only if

$$\frac{1}{N-1}\sum_{j \ne 2} e_j(t) > 1 - \frac{1}{N-1}\frac{c_2 - b_2}{p_2L_2}$$

that is, only if

$$\frac{1}{N-1}\sum_{j \ne 2} e_j(0) + \frac{1}{N-1} > 1 - \frac{1}{N-1}\frac{c_2 - b_2}{p_2L_2} \ ,$$

which reduces to

$$\frac{1}{N-1}\sum_{j > K} e_j(0) + \frac{1}{N-1} > 1 - \frac{1}{N-1}\frac{c_2 - b_2}{p_2L_2} \ . \tag{11}$$

Player 2's switch causes Player 3 to switch too only if

$$\frac{1}{N-1}\sum_{j \ne 3} e_j(t) > 1 - \frac{1}{N-1}\frac{c_3 - b_3}{p_3L_3} \ ,$$

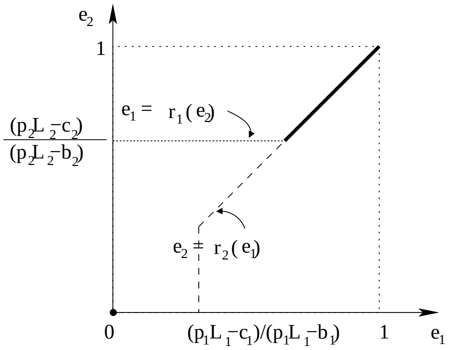

Figure 2: Reaction functions for a two-player weakest-link game. Bold lines and dots indicate potential Nash equilibria.

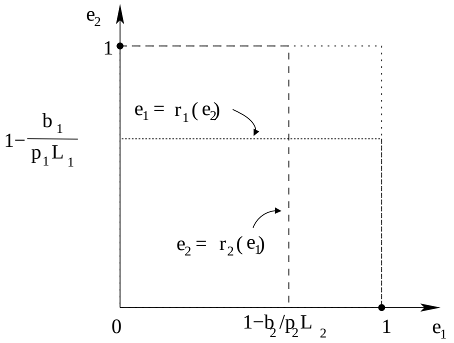

Figure 3: Reaction functions for a two-player best shot game. Bold dots indicate potential Nash equilibria. Protection costs are assumed here to be smaller than insurance costs for both players.

that is,

$$\frac{1}{N-1}\sum_{j>K} e_j(0) + \frac{2}{N-1} > 1 - \frac{1}{N-1}\frac{c_3-b_3}{p_3L_3} . \qquad (12)$$

From Eqs. (11) and (12) we get

$$\frac{c_2-b_2}{p_2L_2} - \frac{c_3-b_3}{p_3L_3} < 1 .$$

Iterating over the K players that are initially not protecting, we get:

$$\max_{2\leq i\leq K}\frac{c_i-b_i}{p_iL_i} - \min_{2\leq i\leq K}\frac{c_i-b_i}{p_iL_i} < K-1 .$$

We can follow an identical derivation for the case where the K players switch from a protection strategy to a non-protection strategy. We then obtain the following necessary condition for "domino effects" to occur over K players, that is a switch in Player 1's strategy causing a switch in the strategy of K players:

$$\left|\max_{2\leq i\leq K}\frac{c_i-b_i}{p_iL_i} - \min_{2\leq i\leq K}\frac{c_i-b_i}{p_iL_i}\right| < K-1 . \qquad (13)$$

Result 2: *We have derived a stability measure of the heterogeneity of a total effort security game with N agents (Eqn. (13)). The more heterogeneous the players are, the more unlikely Eqn. (13) is to hold for large values of K. In other words, the more heterogeneous a system is, the more likely it is to be resilient to perturbations due to a single individual changing strategies.*

4.2 Weakest-link

Here again, we start by considering a two-player game. Computing partial derivatives in e_i and s_i from Eqn. (3), we observe that each player chooses either $(e_i, s_i) = (0, 1)$ (insurance strategy) or $(e_i, s_i) = (\min_{j\neq i} e_j, 0)$ (protection strategy, where in the two-player version of the game $\min_{j\neq i} e_j$ is naturally equal to the protection value chosen by the other player) in order to maximize their utility function.

Looking at the payoffs that can be obtained in both cases leads us to the reaction functions of both players, which we plot in Figure 2. In the figure, we see that a fixed-point is attained when $e_1 = e_2 = 0$ (insurance-only equilibria) and when both e_1 and e_2 are greater than $\max\{(p_1L_1 - c_1)/(p_1L_1 - b_1), (p_2L_2 - c_2)/(p_2L_2 - b_2)\}$.

Result 3: *Generalizing to N players, we obtain the following distinction for the weakest link security game:*

- *Full protection eq.:* If, for all i, $p_iL_i > b_i$, and either 1) $p_iL_i < c_i$, or 2) $p_iL_i \geq c_i$ and $\hat{e}(0)$, the minimum of the security levels initially chosen by all players, satisfies

  $$\hat{e}(0) > \max_{1\leq i\leq N}\{(p_iL_i - c_i)/(p_iL_i - b_i)\} ,$$

  then we have a Nash equilibrium where everyone picks $(\hat{e}(0), 0)$.

- *Multiple eq. without protection:* All players select $e_i = 0$ if the conditions above do not hold. The value of insurance they select depends on their respective valuations. Players for whom insurance is too expensive ($p_iL_i < c_i$) do not insure, with $s_i = 0$, while others choose full insurance, that is $s_i = 1$.

The likelihood of reaching a full protection equilibrium is conditioned by the player which has the largest difference between protection and insurance costs relative to its expected losses. In particular, it only takes one player with an insurance premium smaller than its protection cost ($b_i > c_i$) to make the full protection equilibrium unreachable. Hence, when N grows large, we expect protection equilibria to become more and more infrequently observed.

4.3 Best shot

Looking at the variations of the payload function U_i given in Eqn. (4) as a function of e_i and s_i tells us there are three possibilities for maximizing U_i: a passivity strategy $(0, 0)$, a secure-only strategy $(1, 0)$ and an insure-only strategy $(0, 1)$.

We get $U_i(0, 0) = M_i - p_iL_i(1 - \max\{e_{-i}\})$, $U_i(1, 0) = M_i - b_i$, and $U_i(0, 1) = M_i - c_i$. We immediately notice that $b_i > c_i$ leads Player i to never invest in protection: either the player is passive, or she insures. If, on the other hand $b_i \leq c_i$, then player i chooses a protection strategy over a passivity strategy if and only if (b_i assumed greater than 0) we have $\max\{e_{-i}\} < 1 - b_i/p_iL_i$. We plot the reaction functions, in a two-player case, in Figure 3.

Result 4: *For the two-player best shot security game we can identify the following equilibria:*

- *Protection eq.:* In contrast to the homogeneous case a protection equilibrium does exist. The Nash equilibrium is a free-riding equilibrium where one player protects, and the other does not.

- *Multiple eq. without protection:* If $b_i > c_i$ for all player i individuals will choose to self-insure or remain passive.

In the homogeneous version of the game, we had noted that these Nash equilibria were not reached in a synchronized game with N players, as players would constantly oscillate between free-riding and protecting [19]. With heterogeneous players, however, it is possible to reach a Nash equilibrium. Indeed, if the initial protection levels chosen satisfy $\max\{e_{-i}(0)\} > 1 - b_i/p_iL_i$ for all players *but one*, this last player will be the only one to secure, while everybody else will defect. Note that there should be only one player choosing to secure for a Nash equilibrium to be reached – as soon as at least two players decide to protect, each will defect in the next round hoping to free-ride on the other protecting players. In other words, if there exists a unique i for which the initial constellation of protection levels satisfies

$$\max\{e_{-i}(0)\} < 1 - b_i/p_iL_i \; , \tag{14}$$

then a Nash equilibrium where all players free-ride on player i is reached as long as $b_i < c_i$. This situation could happen when only one player faces disproportionate losses compared to other players, or her security costs are very small.

Result 5: *When protection levels are initially randomly set, protection equilibria in the best shot game are increasingly unlikely to happen as the number of players N grows.*

Assume that the initial protection levels, $e_i(0)$ for $1 \leq i \leq N$ are set independently and at random, that is, that they can be expressed as a random variable with cumulative distribution function F. Then for any Player k, the probability that $e_k(0) < 1 - b_i/p_iL_i$ is simply $F(1 - b_i/p_iL_i)$. It follows that Eqn. (14) is satisfied for Player i with probability $F(1 - b_i/p_iL_i)^{N-1}$.

Next, we want Eqn. (14) to be violated for all players other than i. Eqn. (14) is defeated for a given Player k with probability $1 - F(1 - b_k/p_kL_k)^{N-1}$. Consequently, it is defeated for all Players $j \neq i$ with probability $\prod_{j \neq i}(1 - F(1 - b_j/p_jL_j)^{N-1})$.

It follows that the probability ρ_i that Eqn. (14) is satisfied *only* for Player i is given by

$$\rho_i = F\left(1 - \frac{b_i}{p_iL_i}\right)^{N-1} \prod_{j \neq i}\left(1 - F\left(1 - \frac{b_j}{p_jL_j}\right)^{N-1}\right) \; .$$

Then, the probability that a protection equilibrium can be reached is given by $\sum_i \rho_i$, since the ρ_i's characterize mutually exclusive events. To simplify notations, let $x_i = F\left(1 - \frac{b_i}{p_iL_i}\right)$. Rearranging terms gives

$$\sum_i \rho_i = \sum_i \prod_{j \neq i}\left(1 - x_j^{N-1}\right) - N \prod_j \left(1 - x_j^{N-1}\right) \; .$$

Let $k = \arg\max_i \left\{\prod_{j \neq i}\left(1 - x_j^{N-1}\right)\right\}$. Then we have

$$\sum_i \rho_i \leq N \prod_{j \neq k}\left(1 - x_j^{N-1}\right) - N \prod_j \left(1 - x_j^{N-1}\right) \; ,$$

which gives us, after rearranging

$$\sum_i \rho_i \leq N x_k^{N-1} \prod_{j \neq k}\left(1 - x_j^{N-1}\right) \; ,$$

which tends to zero as N increases, as soon as $x_k = F(1 - b_k/p_kL_k) < 1$.

This is notably the case if we assume a function F strictly monotonous increasing on $[0, 1]$, and positive security costs ($b_i > 0$) for all players.

4.4 Weakest target

As in the homogeneous case [19], Nash equilibria for the weakest target game are quite different depending on whether or not we are considering that mitigation is possible.

Without mitigation. In the weakest target game without mitigation, we have reported [19] that, in the homogeneous case where $b_i = b$, $c_i = c$, $p_i = p$ and $L_i = L$, there are no pure strategy Nash equilibrium. The proof can be extended to the heterogeneous case, as we discuss next.

Let us assume that the minimum protection level over all players is set to $\hat{e} < 1$. Then, we can group players in two categories: those who play $e_i = \hat{e}$, and those who set $e_i > \hat{e}$. By straightforward dominance arguments coming from the description of the payoffs in Eqn. (6), players who select $e_i > \hat{e}$ select $e_i = \hat{e} + \varepsilon$, where $\varepsilon > 0$ is infinitesimally small, and $s_i = 0$. Let

$$\varepsilon < \min_i \left\{\frac{p_iL_i}{2b_i}(1 - s_i) + \frac{c_is_i}{2b_i}\right\} \; .$$

Players who play $e_i = \hat{e}$ would actually prefer to switch to $\hat{e} + 2\varepsilon$. Indeed, the switch in strategies allows a payoff gain of

$$U_i(\hat{e} + 2\varepsilon, 0) - U_i(\hat{e}, s_i) = -2b_i\varepsilon + p_iL_i(1 - s_i) + c_is_i > 0 \; .$$

Hence, this strategy point is not a Nash equilibrium. It follows that the only possible equilibrium point would have to satisfy $e_i = 1$ for all e_i. However, in that case, all players are attacked, which ruins their security investments. All players therefore have an incentive to instead select $e_i = \hat{e} = 0$, which, per the above discussion, cannot characterize a Nash equilibrium.

Result 6: *In the weakest-target game without mitigation we find that pure Nash equilibria for non trivial values of b_i, p_i, L_i and c_i do not exist.*

With mitigation. In the weakest target game with mitigation, we showed that, with homogeneous agents, a full protection Nash equilibrium exists as long as protection costs are smaller than insurance costs [19]. An exactly identical proof can be conducted in the heterogeneous case to show that a full protection equilibrium is reached if $b_i < c_i$ for all i.

On the other hand, it only takes one of the players to face high security costs to make this equilibrium collapse. Indeed, if there exists k such that $b_k > c_k$, then Player k will always prefer a full-insurance strategy ($(e_k, s_k) = (0, 1)$) over a full-protection strategy ($(e_k, s_k) = (1, 0)$). This will immediately lead other players to try to save on security costs by picking $e_i = \varepsilon > 0$ as small as possible. We then observe an escalation as in the unmitigated version discussed above. Hence, heterogeneity actually threatens the (precarious) stability of the only possible Nash equilibrium.

Result 7: In contrast to the weakest-target game without mitigation we find that a pure Nash equilibrium may exist.

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

5. INTERVENTION MECHANISMS

In practice system designers may not be satisfied with the outcomes predicted by non-cooperative game theory. First, equilibria may not be achievable due the complexity of the games, which limits the understanding and accurate execution of strategies by agents. Second, planners may wish to improve upon the Nash equilibrium security practices. Below we discuss selected intervention strategies in the context of the security games to improve convergence and to achieve certain contribution targets.

Objective 1 - Help agents to identify individually rational strategy: In the five games we consider, agents will incur a loss when adequate protection or self-insurance is amiss. However, the reasons for vulnerability to a loss and eventual compromise are different. For example, in the weakest target game without mitigation, a security breach is not solely the result of an agent's protection level, but is dependent on the ordering of contribution levels. Individual rationality presumes that agents follow a sophisticated mixed strategy [19]. However, non-automated agents will only be able to follow such a strategy with difficulty [37]. Even pure strategies might require several periods of convergence [9].

One possible method of intervention to overcome complexity or coordination problems is to offer (non-binding) advice to agents in a security game. For example, Brandts and MacLeod [6] show that players might choose, in a self-enforcing manner, a strategy recommended by an external arbiter. The assignment strongly influences behavior if it does not conflict with another focal principle. In practice, individuals care about who is giving the advice. For example, the suggestion by a computer security company to protect against security breaches with a product of the same brand might be regarded as advertisement and be less influential [27]. Instruments for coordination may also take the form of financial incentives. For example, a third party or intermediary such as an Internet Service Provider (ISP) can offer a rebate or service discount to its subscribers who demonstrably invest in an adequate level of protection.

System designers have also started to exploit individuals' preferences for status quo settings [23]. If users rarely alter default settings, the importance of choosing secure defaults on the two dimensions of self-insurance and protection is immensely high. For example, the Windows XP firewall, when first introduced to the Microsoft Windows operating system in 2001, was disabled by default. Subsequent to the Blaster worm attack, the default setting was changed to "fully enabled" with Windows Server 2003. As another example, OpenBSD's "secure by default" philosophy means that all non-essential services are disabled by default. This promotes general network security and also encourages users to learn more about potential consequences of making changes to security settings.

Objective 2 - Achieve social optimality: In the weakest link security game, deviation of a single agent i from a full protection strategy can render all other agents' efforts meaningless or force them to self-insure or be passive. However, this decision by agent i can be individually rational if $b_i \geq c_i$ or $b_i > p_i L_i$. The traditional solution to this situation has been to involve a social planner who can mandate certain protection and self-insurance settings that optimize overall system utility. In [19], we discussed social optimum outcomes for the homogeneous agents scenario.

A different approach is to allow agents to conduct *binding* pregame communication. For example, consider a scenario in which an agent can propose to another agent that she will only protect if the other agent agrees to reciprocate. Such two-sided communication can increase the protection contribution in the total effort game since the responding agent can internalize the potential contribution of the proposing agent (rather than merely evaluating $b_i < p_i L_i / N$). In a different scenario, an agent may commit to a high or low protection level and not require reciprocation. Given this one-way pregame communication, the protection contributions by other agents will be unaffected in a total effort game, but impacted in the best shot and weakest target games. In the best shot game, a one-way message indicating that the sender will shirk can encourage another agent i to take action (if $b_i < p_i L_i$). In the weakest target game, the same message would signal to other agents that the sender will bear the burden of the attack. This act of altruism is particularly likely if the agent can self-insure at low cost. Finally, binding pregame communication is largely ineffective for the weakest link game. However, it can help to coordinate on the protection Nash equilibrium that Pareto-dominates other equilibria with a lower $\hat{e}$.

Objective 3 - Overcome free-riding and lack of protection in networks: Free-riding occurs in our games at several points. For example, in the best shot game, agents coordinate so that only one agent exercises maximum protection effort in a protection Nash equilibrium. In the social optimum for the weakest target games with homogeneous agents, one node (that may invest in self-insurance) will bear the brunt of an attack while others shirk [19]. Both outcomes, while maximizing utility, might result in loss of camaraderie and willingness to contribute in the future.

One strategy to probabilistically increase contributions by agents is to leverage the *strategic uncertainty* when agents act independently. The coordination problems inherent in the best shot game with heterogeneous agents and in the weakest target games may lead agents to contribute to protection levels above the social optimum. In practice such an approach might have the merit of increasing general preparedness against different types of attacks. Strategic uncertainty is often a function of network size. For example, in the weakest target game, agents in small groups will notice the increased interdependency and risk of being the weakest target. These agents will decide to self-insure their resources more often $((e_i, s_i) = (0, 1))$. That means with increasing network size, we would observe that more individuals contributing to protection. This result stands in contrast to the weakest link game analysis. In the heterogeneous as well as in the homogeneous game, it becomes increasingly unlikely that contributions to protection are made. Heterogeneity can also moderate protection contributions in a different way. In the homogeneous best shot game, we do not observe individually rational protection contributions at all since agents cannot overcome the associated coordination problems. However, they can achieve higher protection levels when agents have heterogeneous tastes.

Contributions can be increased behaviorally by modifying the framing of a security situation. Framing effects occur when two logically equivalent (but not transparently equivalent) statements describing a problem drive individuals to choose dissimilar options. More specifically, such differences in the presentation may draw a subject's attention to alternative aspects of a decision situation, leading an individual to make mistakes in pursuing her underlying preferences [33]. For example, homogeneous agents can be tempted to contribute in a best shot game if they receive feedback that highlights the uniqueness of their contributions [25]. Similarly, increased protection investments may arise if agents perceive a security situation as more threatening. However, underinvestment can result from resignation with respect to the complexity of the security problem.

6. CONCLUSIONS

6.1 Summary

We model security decision-making by heterogeneous agents in a selection of five games. Some of these games have historical foundations in public good theory (weakest-link, best-shot, and total effort) whereas others were proposed recently (weakest target, with or without mitigation). Agents have two security actions at their disposal. They can contribute to a network-side protection pool or invest in a private good to limit losses.

In prior work we have studied homogeneous populations of users, where all participants have the same utility function. In practice, the 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.

However, the fact that the Internet is increasingly used as a common vector between different businesses, and even as a bridge between completely different user bases – for instance, acting as a bridge between mobile phone networks, home users, and e-commerce retailers, emphasizes the need for considering heterogeneous agents, even though the games considered may become far less tractable. In this paper we present an analysis that considers heterogeneous agents. We find several key differences. For example, we found that in the total effort game stability increases with more pronounced heterogeneity in the agent population. The existence of a protection equilibrium in the weakest link game is threatened if only one agent prefers to self-insure or to remain passive. In the best shot game heterogeneous agents can overcome coordination problems more easily, so that a protection equilibrium is now possible, even though reaching this equilibrium grows increasingly unlikely with a larger number of agents participating in the network. Surprisingly, predictions for pure Nash equilibria of the weakest target games remain unchanged. However, mixed strategies do now have to take consideration of the heterogeneity of agents, and are likely to be intractable analytically.

We discuss several intervention strategies in the context of security games. We note that in each game the challenge to increase security contributions to achieve a particular objective requires a largely different approach. This versatility is confirmed by practical observations which tell us that a "one size fits all strategy" for computer security does not exist.

6.2 Future research directions

First, we wish to extend our analysis to more formally explain the impact of limited information on agents strategies. In particular, in computer security and distributed networks the assumption of full information is useful as a first approximation but requires further validation. Similarly, we plan to analyze the robustness of our model by studying the influence of other simplifying assumption (e.g., linear cost parameters). Furthermore, we intend to evaluate strategy changes if moves are conducted sequentially rather than simultaneously. In the context of decision making on the Internet Friedman et al. also distinguish between synchronous and asynchronous moves [16].

Second, 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.

We are determined to incorporate our findings in updated models of system security. In prior work we challenged the assumption that all players are perfectly rational. In [11] we assumed agents to also accept strategies that are near rational and studied how system convergence prediction change.

Our research agenda of formal analysis combined with laboratory experiments is aimed to increase the understanding of individual and organizational security decision making. However, we are also interested in the design of meaningful security policies and aim at developing actionable guidelines for IT managers and other practitioners.

7. ACKNOWLEDGMENTS

We appreciate the detailed evaluations of the anonymous reviewers. This work is supported in part by the National Science Foundation under ITR award ANI-0331659 (100x100) and CCF-0424422 (TRUST - Team for Research in Ubiquitous Secure Technology).

8. 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] R. Anderson. Why information security is hard - an economic perspective. In *Proceedings of the 17th Annual Computer Security Applications Conference (ACSAC'01)*, New Orleans, LA, December 2001.

[3] R. Anderson and T. Moore. The economics of information security. *Science*, 314(5799):610–613, October 1998.

[4] AOL/NSCA. Online safety study, October 2004. Available at: `http://www.security.iia.net.au/downloads/safety_study_v04.pdf`.

[5] T. August and T. Tunca. Network software security and user incentives. *Management Science*, 52(11):1703–1720, November 2006.

[6] J. Brandts and W. MacLeod. Equilibrium selection in experimental games with recommended play. *Games and Economic Behavior*, 11(1):36–63, October 1995.

[7] Bruskin Research. Nearly one in four computer users have lost content to blackouts, viruses and hackers according to new national survey, 2001. Condensed results available at: `http://www.corporate-ir.net/ireye/ir_site.zhtml?ticker=iom&script=410&layout=-6&item_id=163653`.

[8] J.A. Bull, L. Gong, and K. Sollins. Towards security in an open systems federation. In *Proceedings of the Second European Symposium on Research in Computer Security (ESORICS), Springer LNCS No. 648*, pages 3–20, Toulouse, France, November 1992.

[9] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ, 2003.

[10] P. Chen, G. Kataria, and R. Krishnan. On software diversification, correlated failures and risk management, April 2006. Available at SSRN: `http://ssrn.com/abstract=906481`.

[11] N. Christin, J. Grossklags, and J. Chuang. Near rationality and competitive equilibria in networked systems. In *Proceedings of ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, pages 213–219, Portland, OR, August 2004.

[12] D. Clark, J. Wroclawski, K. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow's Internet. In *Proc. of ACM SIGCOMM'02*, pages 347–356, Pittsburgh, PA, Aug. 2002.

[13] G. Danezis and R. Anderson. The economics of resisting censorship. *IEEE Security & Privacy*, 3(1):45–50, January–February 2005.

[14] I. Ehrlich and G.S. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, July 1972.

[15] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of 14th ACM Conference on Computer and Communications Security (CCS)*, pages 375–388, Alexandria, VA, October 2007.

[16] E. Friedman, M. Shor, S. Shenker, and B. Sopher. An experiment on learning with limited information: nonconvergence, experimentation cascades, and the advantage of being slow. *Games and Economic Behavior*, 47(2):325–352, May 2004.

[17] D. Geer, C. Pfleeger, B. Schneier, J. Quarterman, P. Metzger, R. Bace, and P. Gutmann. Cyberinsecurity: The cost of monopoly. How the dominance of Microsoft's products poses a risk to society, 2003. Available from Computer & Communications Industry Association at `http://www.ccianet.org/papers/cyberinsecurity.pdf`.

[18] L.A. Gordon and M. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, November 2002.

[19] J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 2008 World Wide Web Conference (WWW'08)*, pages 209–218, Beijing, China, April 2008.

[20] J. Hartley. Retrospectives: The origins of the representative agent. *The Journal of Economic Perspectives*, 10(2):169–177, Spring 1996.

[21] 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, December 2006.

[22] J. Hirshleifer. From weakest-link to best-shot: the voluntary provision of public goods. *Public Choice*, 41(3):371–386, January 1983.

[23] D. Kahneman and A. Tversky. *Choices, values and frames*. Cambridge University Press, Cambridge, UK, 2000.

[24] S. Kandula, D. Katabi, M. Jacob, and A. Berger. Botz-4-sale: Surviving organized DDoS attacks that mimic flash crowds. In *Proceedings of the 2nd USENIX Symposium on Networked Systems Design & Implementation (NSDI'05)*, pages 287–300, Boston, MA, May 2005.

[25] S. Karau and K. Williams. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65(4):681–706, October 1993.

[26] S. H. Khor, N. Christin, T. Wong, and A. Nakao. Power to the people: Securing the Internet one edge at a time. In *Proceedings of ACM SIGCOMM'07 Workshop on Large-Scale Attack Defense (LSAD)*, pages 89–96, Kyoto, Japan, August 2007.

[27] J. Kuang, R. Weber, and J. Dana. How effective is advice from interested parties?: An experimental test using a pure coordination game. *Journal of Economic Behavior and Organization*, 62(4):591–604, April 2007.

[28] S. Malphrus. The "I Love You" computer virus and the financial services industry, May 2000. Testimony before the Subcommittee on Financial Institutions of the Committee on Banking, Housing, and Urban Affairs, U.S. Senate. `http://www.federalreserve.gov/BoardDocs/testimony/2000/20000518.htm`.

[29] 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.

[30] D. Moore, C. Shannon, and J. Brown. Code-Red: a case study on the spread and victims of an internet worm. In *Proceedings of 2nd ACM/USENIX Internet Measurement Workshop*, pages 273–284, Marseille, France, November 2002.

[31] A. O'Donnell and H. Sethu. On achieving software diversity for improved network security using distributed coloring algorithms. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS)*, pages 121–131, Washington, DC, October 2004.

[32] N. Provos, D. McNamee, P. Mavrommatis, K. Wang, and N. Modadugu. The ghost in the browser: Analysis of web-based malware. In *Proceedings of the First USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)*, Cambridge, MA, April 2007.

[33] M. Rabin. Psychology and economics. *Journal of Economic Literature*, 36(1):11–46, March 1998.

[34] E. Rescorla. Security holes... who cares? In *Proceedings of the 12th USENIX Security Symposium*, pages 75–90, Washington, DC, August 2003.

[35] T. Sandler and K. Hartley. Economics of alliances: The lessons for collective action. *Journal of Economic Literature*, XXXIX(3):869–896, September 2001.

[36] S. Saroiu, S. Gribble, and H. Levy. Measurement and analysis of spyware in a university environment. In *Proceedings of the 1st USENIX Symposium on Networked Systems Design & Implementation (NSDI'04)*, pages 141–153, San Francisco, CA, 2004.

[37] J. Shachat and J.T. Swarthout. Do we detect and exploit mixed strategy play by opponents? *Mathematical Methods of Operations Research*, 59(3):359–373, July 2004.

[38] The Honeynet Project. Know your enemy: the tools and methodologies of the script-kiddie, July 2000. Available online at `http://project.honeynet.org/papers/enemy/`.

[39] H.R. Varian. System reliability and free riding. In L.J. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

[40] N. Weaver, D. Ellis, S. Staniford, and V. Paxson. Worms vs. perimeters: the case for hard-LANs. In *Proceedings of the 12th Annual IEEE Symposium on High Performance Interconnects*, pages 70–76, Stanford, CA, August 2004.

[41] N. Weaver and V. Paxson. A worst-case worm. In *Proceedings (online) of the Third Annual Workshop on Economics and Information Security (WEIS'04)*, Minneapolis, MN, May 2004. Available at `http://www.dtc.umn.edu/weis2004/weaver.pdf`.

[42] L. Zhuang, J. D. Tygar, and R. Dhamija. Injecting heterogeneity through protocol randomization. *International Journal of Network Security*, 4(1):45–58, January 2007.