When Information Improves Information Security ⋆

Jens Grossklags[a], Benjamin Johnson[b], and Nicolas Christin[b]

[a]School of Information, UC Berkeley, [b]CyLab, Carnegie Mellon University
jensg@ischool.berkeley.edu
{johnsonb,nicolasc}@andrew.cmu.edu

The disconnect between meager user investments in security technology and the resulting potential losses can be partially explained by negative externalities: that is, the level of effective protection that security-conscious users obtain is considerably lowered by the insecure behavior from their peers, which in turn provides a personal disincentive to invest in security primitives. Likewise, the lack of accurate understanding of threats is commonly held to significantly weaken the quality of security decision-making.

This paper moves toward a formal, quantitative evaluation of the impact of bounded-rational security decisions in the presence of limited information availability and externalities. We investigate a mixed economy of an individual rational expert and several naïve near-sighted agents in the security decision-making context. We model three canonical types of negative externalities (weakest-link, best shot and total effort), and study the impact of two information regimes on the threat level agents are facing.

We present a methodology to determine and compare strategies and payoffs between the different interdependencies and information conditions. To quantify the agents' valuation of better threat information we propose a metric formalization: the payoff under complete information divided by the payoff under incomplete information. We study this ratio metric analytically and numerically and isolate parameter regions where being more informed creates a payoff advantage for the expert agent.

1 Introduction

Users frequently fail to deploy, or upgrade security technologies, or to carefully preserve and backup their valuable data [26, 35], which leads to considerable monetary losses to both individuals and corporations every year. A partial interpretation of this state of affairs is that *negative externalities* impede end-users' investments in security technologies [43]. Negative network externalities occur when the benefit derived from adopting a technology depend on the actions of others as is frequently the case in the context of network security. For example, users who open and respond to unsolicited advertisements increase the load of spam for all participants in the network, including participants who are making the effort to adopt secure practices. Similarly, choosing a weak password for a corporate VPN system can facilitate compromises of many user accounts, possibly including those of individuals with strong passwords if trust relationships inside the VPN exist.

In other words, a rational user facing negative externalities could make the decision *not* to invest in security primitives given that their personal investment may only marginally matter if other users are adopting insecure practices, or if the perceived cost of a security breach significantly exceeds the cost of investing in security [24].

Moreover, this risk management explanation overemphasizes the rationality of the involved consumers [25]. In practice, consumers face the task to "prevent security breaches within systems that sometimes exceed their level of understanding" [8]. In other words, the amount of information users may be able to acquire and/or to process, is much more limited than is required for a fully rational choice.

⋆ We thank John Chuang for his helpful comments to an earlier version of this paper. All remaining errors are our own. This work is supported in part by the National Science Foundation under ITR award ANI-0331659 (100x100) and with a University of California MICRO project grant in collaboration with DoCoMo USA Labs.

We anticipate the vast majority of users to be *non-expert*, and to apply approximate decision-rules that fail to accurately appreciate the impact of their decisions on others [3]. In particular, in this paper, we assume non-expert users to conduct a simple self-centered cost-benefit analysis, and to neglect externalities. Such users would secure their system only if the vulnerabilities being exploited can cause significant harm or a direct annoyance to them (e.g., their machines become completely unusable), but would not act when they cannot perceive or understand the effects of their insecure behavior (e.g., when their machine is used as a relay to send moderate amounts of spam to third parties). In contrast, an advanced, or expert user fully comprehends to which extent her and others' security choices affect the network as a whole, and responds rationally.

The first contribution of this paper is to study the strategic optimization behavior of such an expert user in an economy of inexperienced users, using three canonical security games that account for externalities [21]. That is, we investigate the extent to which expert users, who understand the nature of the externalities in play, are in a better position to make informed security decisions compared to non-experts. We thus provide a first step toward a formal, quantitative understanding of the impact of negative externalities on security decision-making.

Our approach to capture bounded-rational behaviors of end-users differs significantly from research on computability and approximation of economic equilibria. We argue that models of security decision-making can benefit from a critical inquiry of the conceptual understanding users have of security problems. While experts and unsophisticated users co-exist in the same networks, they do not share the same knowledge or mental models about security problems and countermeasures [8], or the same identical perfectly rational approaches to solve security issues [1, 13].

The second contribution of this paper is to address how the security choices by users are mediated by the information available on the severity of the threats the network faces. We assume that each individual faces a randomly drawn probability of being subject to a direct attack. Indeed in practice, different targets, even if they are part of a same network, are not all equally attractive to an attacker: a computer containing payroll information is, for instance, considerably more valuable than an old "boat anchor" sitting under an intern's desk. Likewise, a machine may be more attractive than another due to looser restrictions in the access policies to the physical facility where the machine is located.

We study how the decisions of the expert and unsophisticated users differ if all draws are common knowledge, compared to a scenario where this information is only privately known – that is, when each player only knows their own probability of being attacked, but not the specific probabilities *other* users may be targeted. We further evaluate the value of better information on the total expected payoff of the expert agent. Specifically, we study the following metric: the payoff under the complete information condition divided by the payoff under the incomplete information condition.

By evaluating the value of information for a range of parameters in different security scenarios, we can determine which configurations can accommodate limited information environments (i.e., when being less informed does not significantly jeopardize an expert user's payoff), as opposed to configurations where expert users and non-expert users achieve similar outcomes due to a lack of available information. This analysis has implications for network designers that want to avoid undesirable hotspots that penalize users for their lack of information about threats. Similarly, ISPs or other intermediaries may take influence on the pricing and availability of security technologies to steer users to less harmful parameter configurations.

We first discuss selected work related to our analytic model (Section 2). In Section 3, we summarize the security games framework and detail our assumptions about agent behaviors and information conditions. We present our methodology and formal analysis in Section 4. We discuss the results and their implications in Section 5, and conclude in Section 6.

2 Related work

In our prior work we have reviewed the research area of the economics of security in depth [21]. In this paper we conduct a decision-theoretic analysis for a sophisticated (expert) agent who interacts with a group of users that follow a simple but reasonable rule-of-thumb strategy.

Our research complements work on market insurance for security and privacy [5, 44]. In particular, several researchers have investigated the impact of correlation of risks and interdependency of agents in networks on the viability of insurance [9, 10, 38]. We structure the remainder of the review of related literature and background information into three selected areas in which we are making a research contribution.

Our work significantly differs from prior decision-theoretic approaches. Gordon and Loeb present a model that highlights the trade-off between perfect and cost-effective security [19]. They consider the protection of an information set that has an associated loss if compromised, probability of attack, and probability that attack is successful. They show that an optimizing firm will not always defend highly vulnerable data, and only invest a fraction of the expected loss. Cavusoglu *et al.* [12] consider the decision-making problem of a firm when attack probabilities are externally given compared to a scenario when the attacker is explicitly modeled as a strategic player in a game-theoretic framework. Their model shows that if the firm assumes that the attacker strategically responds then in most considered cases its profit will increase. Schechter and Smith [39] consider the decision-theoretic analysis from the perspective of the potential intruder. They highlight several modeling alternatives for attacker behavior and their payoff consequences. The analytic work on security investments and level of penalties for offenses is complemented by empirical research [37, 41].

2.1 Bounded rationality

Acquisti and Grossklags summarize work in the area of behavioral economics and psychology that is of relevance for privacy and security decision-making [3]. Users' decisions are not only limited by cognitive and computational restrictions (i.e., bounded rationality), but are also influenced by systematic psychological deviations from rationality.

Recent research has investigated agents that overemphasize earlier costs and benefits at the expense of their future well-being [1, 4]. Christin *et al.* suggest that agents respond near-rationally to the complexity of networked systems [13]. In their model individuals are satisfied with a payoff within a small margin of the optimal outcome.

Different from the above work that considers all players to act the same, the current paper studies a mixed economy, with expert and non-expert users co-existing. While expert users are as rational as possible, non-expert users deviate from rationality by adopting approximate (rules-of-thumb) decision strategies. In practice, users frequently have to rely on rules-of-thumb when a "quantitative method to measure security levels" is not available [31]. Economic analysis including rule-of-thumb choices have been discussed outside of the security context, e.g., [15] [16] [29].

2.2 Limited information

In the context of the value of security information, research has been mostly concerned with incentives for sharing and disclosure. Several models investigate under which conditions organizations are willing to contribute to an information pool about security breaches and investments when competitive effects may result from this cooperation [18, 20]. Empirical papers explore the impact of mandated disclosures [11] or publication of software vulnerabilities [42] on the financial market value of corporations. Other contributions to the security field include computation of Bayesian Nash outcomes for an intrusion detection game [30], and security patrol versus robber avoidance scenarios [36].

We conduct a comparative analysis of strategies and payoffs for a sophisticated agent in a security model when the likelihood of a directed attack is either common or private knowledge. In particular, we evaluate the influence of the lack of information given different organizational dependencies [43].

2.3 Heterogeneous agents

In our previous work we have analyzed both the case of homogeneous [21] and heterogeneous agents [22]. When considering heterogeneous agents, however, we have focused on differences in the costs agents may face. We assumed that users differ in the price they have to pay for protection and self-insurance, and that they have different perceived or actual losses associated with successful (uninsured) security compromises. In the present paper we analyze the case of agents facing different attack probabilities, that may be a priori unknown to other agents.

Given certain differences in the attractiveness of a particular target the question remains how a defender is able to determine a reasonable estimate of the attack probability. Such a problem far exceeds the scope of this paper, whose main goal is to study the impact of information (or lack thereof) on security strategies, and we refer the reader to the threat modeling literature. (See [6] for an introduction and references.)

3 Decision Theoretic Model

We next summarize the security games we analyze, and extend models we previously proposed [21] to the case of an economy consisting of an expert user and several unsophisticated users.

3.1 Basic model

Self-protection and self-insurance. In practice, the action portfolio of a defender may include different options to prevent successful compromises and to limit losses that result from a breach. In Grossklags *et al.* [21] we provide a model that allows a decoupling of investments in the context of computer security. On the one hand, the perimeter can be strengthened with a higher self-protection investment (e.g., implementing or updating a firewall). On the other hand, the amount of losses can be reduced by introducing self-insurance technologies and practices (e.g., backup provisions). Formally, player i decides whether to invest in protection ($e_i = 1$) or not ($e_i = 0$). Similarly, each player can adopt a self-insurance technology ($s_i = 1$) or not ($s_i = 0$). In other words, e_i and s_i are two discrete decision variables.

Discrete choice decision-making captures many practical security problems. Examples include purchase and adoption investments as well as updating and patching of protection and self-insurance technologies [7, 28, 32, 34].

We have further conducted a sensitivity analysis with respect to the discrete choice assumption and find that, for the study in the present paper, the only differences between the discrete and continuous cases (where e_i and s_i are continuous variables over the interval $(0, 1)$ as opposed to be mere binary variables) arise when there is strict equality between some of the terms in our case-specifying inequality conditions (see derivations in Section 4). We believe that focusing on these boundary cases is of limited practical applicability, and could even be misleading. For comparison, we refer to our prior work where we considered the continuous case in a full information environment [21].

We further denote by $b \geq 0$ and $c \geq 0$ the cost of protection and self-insurance, respectively, which are homogeneous for the agent population. So, player i pays be_i for protection and cs_i for self-insurance.

Interdependency. Decisions by one defender frequently influence the incentives for security investments by her peers [2, 28, 33, 43]. For example, the lack of protection effort by a subset of agents will often allow an attacker to also compromise resources of other agents if a common perimeter is breached. We focus in

this work on such tightly coupled networks [43].[1] In a tightly coupled network all defenders will face a loss if the condition of a security breach is fulfilled whereas in a loosely coupled network consequences may differ for network participants [17]. We denote H as a "contribution" function that characterizes the effect of e_i on agent's utility U_i, subject to the protection levels chosen (contributed) by *all* other players. We require that H be defined for all values over $(0, 1)^N$. We distinguish three canonical cases that we discussed in-depth in prior work [21]. (Following common notation, e_{-i} denotes the set of protection levels chosen by players other than i in the equations below.)

Weakest-link (e.g., network access control): $H = \min(e_i, e_{-i})$.
In the weakest-link game, the network participant with the least amount of protection defines the protection level of the whole network. Weakest-link instances occur quite frequently in practice. For example, most corporate networks are designed to be isolated from the Internet. However, in a vast majority of cases, these corporate networks have a few hosts (e.g., firewall, web server, ...) that have access to both the "public" Internet as well as the "private" networks. For an attacker interested in penetrating the corporate network, compromising only one of these hosts connected to the public Internet suffices to gain access to the private network.

Best shot (e.g., censorship resilience): $H = \max(e_i, e_{-i})$.
In the best shot game, the protection level of the whole network is equal to the protection level of the most protected player. A practical instance of this game is in censorship-resilient systems, where a number of end-hosts are sharing a copy of a document. To make this document unavailable to anybody joining the network, the censor (attacker) has to remove *all* copies of the document from the network. As such, the protection level of the most protected host having a copy available conditions the protection level of the entire network.

Total effort (e.g., parallel file transfers): $H = \frac{1}{N} \sum_k e_k$.
The total effort game models a situation where the level of protection of the entire network depends on the sum of all the network participants' efforts. Parallel ("swarming") file transfers in a peer-to-peer network implement a form of total effort game, when the attacker is primarily interested in degrading the download speeds users experience. Indeed, the more peers in the swarm the attacker takes down, the more the overall file transfer speed decreases. As such, the overall level of protection of the whole network is dependent on the number of participants to the network. Note that, if the attacker wants to make the transferred file completely unavailable, the game then becomes a best-shot game as discussed above.

Attack probabilities, network size and endowment. Each of $N \in \mathbb{N}$ agents receives an endowment M. If she is attacked and compromised successfully she faces a loss L. We assume that each agent i draws an individual attack probability p_i ($0 \leq p_i \leq 1$) from a uniform random distribution.[2] This models the heterogeneous preferences that attackers have for different targets, due to their economic, political, or reputational agenda. The choice of a uniform distribution ensures the analysis remains tractable, while already providing numerous insights. We conjecture that different distributions (e.g., power law) may also be appropriate in practice.

[1] There is an ongoing debate whether researchers should assume full connectivity of a network graph given modern computer security threats such as worms and viruses. (Personal communication with Nicholas Weaver, ICSI.)

[2] Technically, our analysis does not require complete knowledge of the distribution on attack probabilities. The distribution informs the probability that a given number of p_j are above the rule-of-thumb threshold; but to conduct our analysis, it suffices to know only these threshold probabilities, and not the full distribution.

3.2 Player behavior

At the core of our analysis is the observation that expert and non-expert users differ in their understanding of the complexity of networked systems. Indeed, consumers' knowledge about risks and means of protection with respect to privacy and security can be quite varied [3], and field surveys separate between high and low expertise users [40].

Sophisticated (expert) user. Advanced users can rely on their superior technical and structural understanding of computer security threats and defense mechanisms, to analyze and respond to changes in the environment [14]. In the present context, expert users, for example, have less difficulty to conclude that the goal to avoid censorship points at a best shot scenario, whereas the protection of a corporate network frequently suggests a weakest-link situation [21]. Accordingly, a sophisticated user correctly understands her utility to be dependent on the interdependencies that exist in the network:

$$U_i = M - p_i L(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i \ .$$

Naïve (non-expert) user. Average users underappreciate the interdependency of network security goals and threats [3] [40]. We model the *perceived* utility of each naïve agent to only depend on the direct security threat and the individual investment in self-protection and self-insurance. The investment levels of other players are *not* considered in the naïve user's decision making, despite the existence of interdependencies. We define the perceived utility for a specific naïve agent j as:

$$PU_j = M - p_j L(1 - s_j)(1 - e_j) - be_j - cs_j \ .$$

Clearly, perceived and realized utility actually differ: by failing to incorporate the interdependencies of all agents' investment levels in their analysis, naïve users may achieve sub-optimal payoffs that actually are far below their own expectations. This paper does not aim to resolve this conflict, and, in fact, there is little evidence that users will learn the complexity of network security over time or are able to keep up with the challenges of novel threats [40]. We argue that non-expert users would repeatedly act in an inconsistent fashion. This hypothesis is supported by findings in behavioral economics that consumers repeatedly deviate from rationality, however, in the same predictable ways [27].

3.3 Information conditions

Our analysis is focused on the decision making of the expert user subject to the bounded rational behaviors of the naïve network participants. That is, more precisely, the expert agent maximizes her expected utility subject to the available information about other agents' drawn threat probabilities and their resulting actions. Two different information conditions may be available to the expert agent:

Complete information: Actual draws of attack probabilities p_j for all $j \neq i$, and her own drawn probability of being attacked p_i.

Incomplete information: Known probability distribution of the naïve users' attack threat, and her own drawn probability of being attacked p_i.

The expert agent can accurately infer what each agent's investment levels are in the complete information scenario. Under incomplete information the sophisticated user has to develop an expectation about the actions of the naïve users.

4 Analysis methodology

In the remainder of this discussion, we will always use the index i to denote the expert player, and $j \neq i$ to denote the naïve players. For each of the three games, weakest-link, best shot, and total effort, our analysis proceeds via the following five-step procedure.

1. Determine player i's payoff within the game for selected strategies of passivity, full self-insurance, and full protection. As shown through a relatively simple analysis [21] [22], player i can maximize her utility only by relying on (one or more of) these three strategies.
2. Determine the conditions on the game's parameters (b, c, L, N, p_i, and if applicable, p_j for $j \neq i$) under which player i should select each strategy.
3. Determine additional conditions on the game's parameters such that the probability (relative to p_i) of each case, as well as the expected value of p_i within each case can be easily computed.
4. Determine player i's total expected payoff relative to the distribution on p_i and all other known parameters.
5. In the case of complete information, eliminate dependence on p_j for $j \neq i$ by taking, within each parameter case, an appropriate expected value.

Diligent application of this method generates a table recording the total expected payoffs for player i, given any valid assignment to the parameters b, c, L, N. In the process it also generates tables of selection conditions, probabilities, and expected payoffs for each possible strategy; and in the complete information case, gives results for total expected payoffs conditioned on the exact draws of p_j by the other players. The results are presented in Tables 1–9, located in the appendix.

Due to space limitations, tables involving probabilities and expected payoffs for various strategies as well as intermediate expected total payoff tables conditioned on other players have been omitted, but are available in a companion technical report [23], along with related derivations.

In the remainder of this section we illustrate this method by considering, for each step listed above, one game and one parameter case for which we have applied the appropriate step.

Step 1 example: Passivity payoff computation. Let us consider the challenge of determining payoffs for player i's passivity in the best shot game under the conditions of limited information and parameter constraints $b \leq c$. The general payoff function for the best shot game is obtained by substituting $H(e_i, e_{-i}) = \max(e_i, e_{-i})$ into the general utility function for all games, i.e. $U(i) = M - p_i L(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i$. Doing this, we obtain $U(i) = M - p_i L(1 - s_i)(1 - \max(e_i, e_{-i}) - be_i - cs_i$. To get the payoff for player i's passivity we plug in $e_i = s_i = 0$ to obtain

$$U_i = \begin{cases} M - p_i L, \text{ if } \max_{j \neq i} e_j = 0 \\ M, \text{ if } \max_{j \neq i} e_j = 1 \end{cases} .$$

Now in the incomplete information case, we do not know any of the p_j for $j \neq i$, so we do not know all the parameters to compute the required payoff. However, since we assume that the p_j are independently and uniformly distributed in $[0, 1]$, we can compute an expected value for this payoff as follows. The probability (over p_j) that none of the other players protects (i.e. that $\max_{j \neq i} p_j < b/L$) is exactly $(b/L)^{N-1}$, and in this case the payoff would be $M - p_i L$. The probability (over p_i) that at least one of the other players protects (i.e. that $b/L \leq \max_{j \neq i} p_j$) is exactly $1 - (b/L)^{N-1}$, and in this case the payoff would be M. Thus, the total expected payoff for selecting the passivity strategy is $(b/L)^{N-1} \cdot (M - p_i L) + (1 - (b/L)^{N-1}) \cdot M$, which simplifies to $M - p_i L(b/L)^{N-1}$. We record this as the payoff result for passivity in the incomplete game, with $b \leq c$, as can be seen in Table 4.

Step 2 example: Strategy selection. Let us next consider the challenge of determining parameter conditions under which we should select player i's strategy in the weakest link game. Since this is a second step,

consider the game payoffs in Table 1 as given. We are interested in determining player i's strategic play for any given parameter case. Select for consideration the case $b \leq c$ with incomplete information. (Note that this is the most difficult case for this game).

To determine the optimal strategy for player i, we must select the maximum of the quantities Passivity: $M - p_i L$, Self-insurance: $M - c$, and Protection: $M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose passivity if it is better than self-insurance or protection, i.e. $M - p_i L > M - c$ and $M - p_i L > M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose self-insurance if it is better than passivity or protection, i.e. $M - c \geq M - p_i L$ and $M - c > M - b - p_i L(1 - (1 - b/L)^{N-1})$. We should choose protection if it is better than passivity or self-insurance, i.e. $M - b - p_i L(1 - (1 - b/L)^{N-1}) \geq M - p_i L$ and $M - b - p_i L(1 - (1 - b/L)^{N-1}) \geq M - c$.

Re-writing the above inequalities as linear constraints on p_i, we choose passivity if $p_i \leq c/L$ and $p_i \leq \frac{b}{L(1-(1-b/L)^{N-1})}$; we choose self-insurance if $p_i > c/L$ and $p_i > \frac{c-b}{L(1-(1-b/L)^{N-1})}$; and we choose protection if $\frac{c-b}{L(1-(1-b/L)^{N-1})} \leq p_i \leq \frac{b}{L(1-(1-b/L)^{N-1})}$.

For simplicity of computation, we would like to have our decision mechanism involve only a single inequality constraint on p_i. To obtain this, it is necessary and sufficient to determine the ordering of the three terms: $\frac{c}{L}$, $\frac{b}{L(1-(1-b/L)^{N-1})}$, and $\frac{c-b}{L(1-(1-b/L)^{N-1})}$.

It turns out that there are only two possible orderings for these three terms. The single inequality $c < \frac{b}{(1-b/L)^{N-1}}$ determines the ordering: $\frac{c}{L} < \frac{c-b}{L(1-(1-b/L)^{N-1})} < \frac{b}{L(1-(1-b/L)^{N-1})}$; while the reverse inequality $\frac{b}{(1-b/L)^{N-1}} \leq c$ determines the reverse ordering on all three terms. This observation suggests we should add sub-cases under $b \leq c$ depending on which of these two inequalities holds. See Table 2.

Within each new sub-case the criterion for selecting the strategy that gives the highest payoff can now be represented by a single linear inequality on p_i. If $c \leq \frac{b}{(1-b/L)^{N-1}}$, then passivity wins so long as $p_i < c/L$; (because the new case conditions also guarantee $p_i < \frac{b}{L(1-b/L)^{N-1}}$). Similarly, self-insurance is preferable if $p_i \geq c/L$. Protection never wins in this case because we cannot have $\frac{c-b}{L\left(1-(1-b/L)^{N-1}\right)} \leq p_i \leq \frac{b}{L\left(1-(1-b/L)^{N-1}\right)}$ when we also have $\frac{b}{(1-b/L)^{N-1}} < \frac{c-b}{L\left(1-(1-b/L)^{N-1}\right)}$. The computations for the case $\frac{b}{(1-b/L)^{N-1}} < c$ are similar; the results are recorded in Table 2.

Step 3 example: Case determination. Now, consider the challenge of determining additional constraints on parameters in the total effort game, so that in any given case the total payoffs can be represented by simple closed-form functions of the game's parameters. Since this is a third step, we assume the second step has been diligently carried out and consider the strategy conditions presented in Table 8 as given. For brevity, we consider only the incomplete information case under the assumption $b \leq c$.

To illustrate the problem we are about to face, consider the condition for selecting passivity in the incomplete information game and the case: $b + b^2(N-1)/L < c$. The condition here is that $p_i < bN/L$. This condition is possible if and only if $bN < L$. The case conditions determined thus far do not specify which of these is the case; so for subsequent computations, we will need to know which it is, and therefore must consider the two cases separately.

Going beyond this particular example, there are several other values in this table where a similar phenomenon occurs. In particular, we need new cases to determine whether each of the following relations holds: $bN/L \leq 1$, $\frac{c}{b+(L-b)/N} \leq 1$, and $\frac{c-b}{b-b/N} \leq 1$. (See Table 8). To combine these with previous cases in a way that avoids redundancy, we rewrite the conditions involving c as linear inequalities on c; obtaining $c \leq b + (L-b)/N$ and $c \leq 2b - b/N$.

We are thus left to reconcile these additional cases with the current cases $b \leq c \leq b + \frac{b^2}{L}(N-1)$ and $b + \frac{b^2}{L}(N-1) < c$. To do this efficiently we must know the order of the terms $b + \frac{L-b}{N}$, $2b - \frac{b}{N}$, and $b + \frac{b^2}{L}(N-1)$. Fortunately, it turns out that there are only two possible orderings on these terms; and

furthermore, which of the two orderings it is depends on the relation $bN < L$ which we already needed to specify as part of our case distinctions. If $bN \leq L$, then $b + \frac{b^2}{L}(N-1) \leq 2b - \frac{b}{N} \leq b + \frac{L-b}{N}$ and if $bL > N$, then the reverse relations hold.

Assuming limited information, $b \leq c$, and dividing all cases according to $bN \leq L$, it requires a total of 5 cases to determine all important relationships among important parameters for this game. We may have $bN \leq L$ and $b \leq c \leq b + \frac{b^2}{L}(N-1)$; $bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$; $bN \leq L$ and$2b - \frac{b}{N} \leq c$; $bN > L$ and $c \leq b + \frac{L-b}{N}$; and $bN > L$ and $b + \frac{L-b}{N} < c$. For reference, see Table 9.

Step 4 example: Total payoff computation. Let us determine the total expected payoff for the expert player with incomplete information in the best shot game with $b \leq c$. As intermediate steps we must compute the probability that each strategy is played, along with the expected payoff for each strategy. The total payoff is then given by (Probability of passivity · Expected payoff for passivity) + (Probability of self-insurance · Expected payoff for self-insurance) + (Probability of protection · Expected payoff for protection).

The expected probability of passivity in this case is 1, with a payoff of $M - p_i L(b/L)^{N-1}$. To get an expected payoff, we compute the expected value of p_i within this case. Since there is no constraint on p_i and it is drawn from a uniform distribution its expected value is $1/2$. Thus the expected payoff for this case is $M - (L/2)(b/L)^{N-1}$. The total expected payoff is thus $M - (L/2)(b/L)^{N-1}$.

Step 5 example: Eliminating dependencies on other players. Consider the challenge of examining the total expected payoff for player i, who has complete information, and rewriting this payoff in a way that is still meaningful as an expected payoff, but does not depend on any p_j for $j \neq i$. The reason we want to do this last step is so we can compare complete information payoff results with incomplete information payoff results. We can only do this if the direct dependence on privileged information is removed from the complete information case payoff. Our method of information removal involves taking an appropriate expected value.

For this example, we consider the best shot game with complete information in the case $b \leq c$. Since this is a fifth step, we should assume that the fourth step (i.e., computing the expected payoff for player i as a function of parameters that may include p_j for $j \neq i$) has been accomplished.

Indeed, by following Steps 1–4, the total expected payoffs for player i (conditioned on other players) in the case $b \leq c$ can be derived, subject to two additional sub-cases. If $\max_{j \neq i} p_j \leq b/L$, then the expected payoff is $M - c + c^2/L$; while if $b/L < \max_{j \neq i} p_j$, then the expected payoff is $M - b + b^2/L$.

To generate an appropriate "a posteriori" expected payoff over all choices of p_j, we compute the probability (over choice of p_j) that we are in case $\max_{j \neq i} p_j \leq b/L$ times the payoff for that case, plus the probability (over p_j) that we are in the case $b/L < \max_{j \neq i} p_j$ times the payoff for that case. We obtain $(b/L)^{N-1} \cdot \left[M - c + c^2/L\right] + \left[1 - (b/L)^{N-1}\right] \cdot \left[M - b + b^2/L\right]$. The end result is $M - b\,(1 - b/2L)\,(b/L)^{N-1}$. See Table 6.

5 Results

5.1 Strategies and payoffs

Our results provide us with insights into security decision-making in networked systems. We can recognize several situations that immediately relate to practical risk choices. We start with basic observations that are relevant for all three games, before discussing the different games and information conditions in more detail.

General observations applicable to all three security games. Every scenario involves simple cost-benefit analyses for both sophisticated and naïve agents. Agents remain passive when the cost of self-protection and self-insurance exceeds the expected loss. Further, they differentiate between the two types of security actions based on their relative cost. This behavior describes what we would usually consider as basic risk-taking that is part of everyday life: It is not always worth protecting against known risks.

One important feature of our model is the availability of self-insurance. If $c < b$ the decision scenario significantly simplifies for all games and both information conditions. This is because once self-insurance is applied, the risk and interdependency among the players is removed. The interesting cases for all three games arise when $b \leq c$ and protection is a potentially cost-effective option. In this case self-insurance has a more subtle effect on the payoffs.

There are important differences between the two agent types. The expert agent considers the strategic interdependencies of all agents' choices. For example, consider $b < p_i L$ and $b \leq c$ (that is, protection would be the preferred choice in the absence of interdependencies) then the expert agent sometimes rather prefers to self-insure, or to remain passive while naïve agents would always protect without further consideration. The more nuanced strategies of the expert agent attest to her realization that the group protection goal is sometimes not achievable. Note that we model the agents' incentives to invest in protection in canonical scenarios when security is critically dependent on a group effort (see examples in Section 3). For example, with full cooperation of all agents the incentives to send unsolicited bulk email could be significantly reduced. However, if naïve users open, respond, or otherwise interact with spam then other users have little choice but some form of mitigation of the resulting inconveniences. Otherwise, the expert agent will commonly invest in security for a resource when its safety is not subject to peers' (in)actions (i.e., if $N = 1$).

If $b > p_j L$ for some agents j, then the naïve users do not fully internalize how the inactions of those agents can impact system-wide security. This naïveté is coming back to haunt them. In fact, surveys of average end users' security experiences show that 66 percent lost data permanently due to lacking backup provisions [26]. Similarly, 54 percent have had their computers infected by a network-propagated malicious code [35]. For example, the success of the Conflickr/Downadup worm is dependent on users not applying available patches to their operating system.

The naïve agents face a payoff reduction as a result of their limited understanding of correlated threats, but even the sophisticated agent can experience a similar payoff reduction due to limited information. On the one hand, she might invest in self-protection or self-insurance when it is not necessary because the naïve agents collectively or individually secured the network. On the other hand, she may fail to take a security action when a (relatively unexpected low probability) breach actually occurs. It is important to mention that she acted rationally in both situations, but these additional risks remain.

Basic payoffs for different security actions: We can immediately observe that the additional risk due to limited information results from different mechanisms for each security scenario. In the weakest-link game (Online appendix: Table 1 [**?**]) we find that self-protection carries a risk for the expert agent with limited information that at least one naïve agent chooses not to protect. This would result in a break-down of system security and a waste of self-protection expenditure. In contrast, in the best shot game (Table 4) the investment in preventive action always secures the network but with limited information this may be a duplicative effort. In the total effort game these risks are more balanced (Online appendix: Table 7). The expert can add or withhold her N-th part of the total feasible security contribution. Depending on the cost of security she has to estimate the expected number of naïve contributors K in order to respond adequately.

Conditions for choice between different security actions: In the weakest-link game and complete information, the expert agent can utilize the lowest attack probability that any naïve agent has drawn. If this value is below the required threshold for protection, (i.e. if $\min_{j \neq i} p_j < b/L$), then the sophisticated agent will never protect. Otherwise, depending on her own draw she will make or break a successful defense. Under incomplete information she has to consider the likelihood $(1 - b/L)^{N-1}$ that all naïve agents protect. In all cases there is now a residual likelihood that she might self-insure (Online appendix: Table 2).

In the best shot game the fully informed expert can simply determine the highest likelihood of being attacked for any naïve agent to decide whether she should contribute to system protection. With full or limited information, it is obvious that she will only have to contribute very rarely, and can mostly rely on others' efforts. Nevertheless, it is surprising to find that in the incomplete information scenario the expected

payoff from passivity always dominates the expected payoff for protection, even when the expected loss is near total ($p_i \sim 1$). The sophisticated user with limited information will never protect. Under neither information condition is it optimal to self-insure if $b \leq c$. See Table 5 for details.

Next consider the total effort game (Online appendix: Table 8). Under full information with $b \leq c$, all conditions depend non-trivially on K, the number of contributors to protection. Under incomplete information the expert must compute the expected value of K, which is $(1 - b/L)(N - 1)$. The case differences between complete and incomplete conditions reflect the replacement of K with $E[K]$, and subsequent simplification. In all cases, the critical factor for the decision to protect is whether the potential loss is N times greater than the cost of protection (i.e. $p_i L \geq bN$).

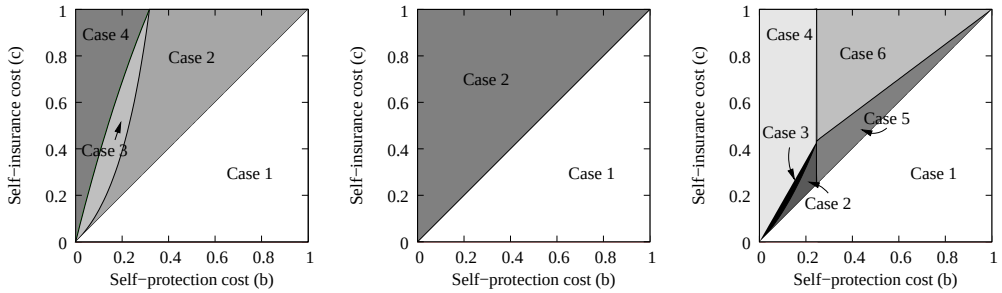

(a) Weakest link (b) Best shot (c) Total effort

Fig. 1. Strategy boundaries in the incomplete information scenario for the sophisticated player. The cases refer to b) Table 6 and a,c) Online app. Tables 3/9 (with $L = M = 1$ and $N = 4$).

Case boundaries for choice between different security actions: In Figure 1, we plot the cases used to record total expected payoffs for the expert agent.

In the weakest-link game only cases 3 and 4 allow for investments in self-protection. We find that increasing the number of agents, N, results in a shrinkage of both cases 3 and 4 to the benefit of case 2. In contrast, the determination of case boundaries in the best shot game is independent of the size of the network. Finally, in the total effort game only cases 3 and 4 allow for rational self-protection investments. Again an increase in the network size reduces the prevalence of these cases (since $bN \leq L$ is a necessary condition).

Payoffs: Table 6, and online appendix Tables 3 & 9 contain the total expected payoff for decisions made by the sophisticated agent, but also for the naïve agents.

We have already highlighted that for $c < b$ all agents follow the same simple decision rule to decide between passivity and self-insurance. Therefore, payoffs in this region are identical for all agent types in the case of homogeneous security costs. But, there are payoff differences among all three information conditions for some parts of the parameter range when $b \leq c$.

Consider the graphs in Figure 2. We plot the payoff functions for sophisticated agents types under the different information conditions, as well as the payoff output for the non-expert agent. It is intuitive that the naïve agents suffer in the weakest-link game since they do not appreciate the difficulty to achieve system-wide protection. Similarly, in the best shot game too many unsophisticated agents will invest in protection lowering the average payoff. In the total effort game, sophisticated agents realize that their contribution is only valued in relation to the network size. In comparison, naïve agents invest more often in protection. This reflects the fact that the naïve agent ignores the self-insurance option whenever protection is cheaper.

We can observe that the sophisticated agents will suffer from their misallocation of resources in the weakest-link game when information is incomplete. In the best shot game this impact is limited, but there is

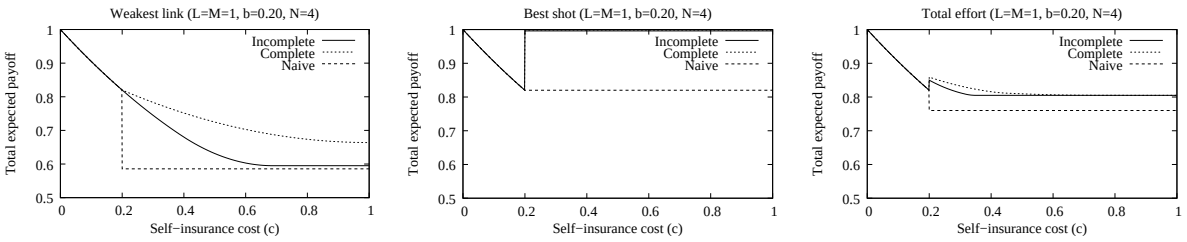

(a) Weakest link (b) Best shot (c) Total effort

Fig. 2. Total expected payoffs for the strategic player under different information conditions, compared with that of the naïve agents (with $L = M = 1$, $N = 4$, and b is fixed to $b = 0.20$).

a residual risk that no naïve agent willingly protects due to an unlikely set of draws. In such cases the fully informed expert could have chosen to take it upon herself to secure the network. In the total effort game we observe a limited payoff discrepancy for expert users as a result of limited information.

5.2 Value of information

From a system design perspective it is important to select parameter settings (e.g., making available specific security technologies) that maximize user utility and are robust to changes in the environment. The security games we analyze in this paper are a significant challenge in both aspects. In particular, from Figure 2 we can infer that the penalty for the lack of complete information about attack threats can be highly variable depending on the system parameters. We argue that the reduction of this disparity should be an important design goal. To further this goal we propose a mathematical formulation to measure the value of better information. We then apply this metric to the analysis of the three canonical security games.

Definition: We are interested in a mathematical measure that allows us to quantify the payoff loss due to incomplete information for sophisticated agents, that can be applied to a variety of decision-theoretic scenarios. It is nontrivial to arrive at a definitive answer for this problem statement, therefore, we consider our analysis as a first step towards this goal. We define the value of information metric as the ratio:

$$\frac{\text{Expected payoff in the complete information environment}}{\text{Expected payoff in the incomplete information environment}}$$

Observations: Consider Figure 3 which gives, for all three security games, a heat plot for the value of better information over all choices of b and c with L, M, N fixed at $L = M = 1$ and $N = 4$. The most remarkable feature of these graphs are the different hotspot regions. In the weakest-link game we find that higher ratios are located within the boundaries of cases 3 and 4. Both cases allow for self-protection in the presence of incomplete information and therefore cover risks more directly than the remaining cases. (Case 1 and 2 associate zero probability with self-protection.)

In the best shot scenario the peak region is located trivially within the boundaries of case 2. We know that the expert player will never protect under incomplete information but is subject to the residual risk of a system-wide security failure. For $N = 4$ the likelihood of such a breakdown is already very small, and decreases with N. Still this outcome is feasible and most pronounced for protection costs that are about a half to two-thirds of the loss, L. For higher b the disincentive of buying self-protection and the potential loss are more balanced resulting in a lower penalty for limited information.

In the total effort game we observe multiple hotspot regions. Cases 4 and 6 are unaffected by limited information. They are characterized by the absence of self-insurance as a feasible strategy. This eases the decision-making problem of the expert, and reduces the likelihood of a misspent security investment.

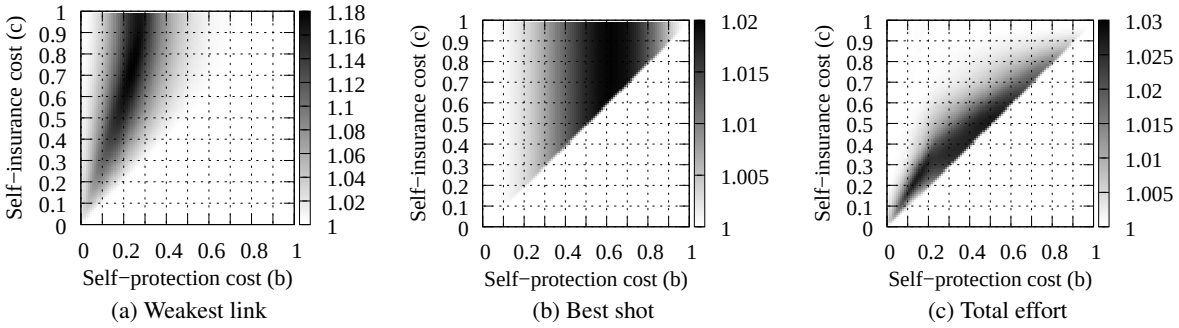

(a) Weakest link (b) Best shot (c) Total effort

Fig. 3. The value of information for the various games (with $L = M = 1$, $N = 4$).

6 Conclusions

In our work we emphasize that security decision-making is shaped by the structure of the task environment as well as the knowledge and computational capabilities of the agents. In our model, decisions are made from three distinct security actions (self-protection, self-insurance or passivity) to confront the security risks of weakest-link, best shot and total effort interdependencies [21, 43]. In these environments, we investigate the co-habitation of a single fully rational expert and $N-1$ naïve agents. The naïve agents fail to account for the decisions of other agents, and instead follow a simple but reasonable self-centered rule-of-thumb. We further study the impact of limited information on the rational agent's choices, and provide a detailed overview and examples of our methodology to compare strategies and payoffs.

We find that in general, the naïve agents match the payoff of the expert when self-insurance is cheap, but not otherwise. Even with limited information, the sophisticated agent can generally translate her better structural understanding into decisions that minimize wasted protection investments, or an earlier retreat to the self-insurance strategy when system-wide security is (likely) failing.

A notable exception is the weakest link game with incomplete information, where the payoff of the sophisticated agent degrades to that of the naïve agent as self-insurance becomes more expensive. A practical implication of this result is that, in corporate network access control, having a lot of information about the various potential vulnerabilities that may exist at network access points actually only marginally enhances security; the key factor is whether self-insurance (e.g., data backups) provide adequate security or not. When some items, such as trade secrets, cannot be self-insured, they simply should not be stored on a publicly accessible network. Common sense tells us that much; a contribution of this paper is to provide a mathematical foundation to justify such policy recommendations.

Our analysis also shows that an expert user never provides a positive improvement to system-wide security (in comparison to her replacement by an unsophisticated agent). While our expert agent is rational, she is not benevolent. Instead she acts selfishly, and the set of scenarios for which protection is her best option is always a subset of the set of scenarios for which the naïve agent chooses protection. In other words, assuming that competent CISOs may be interested in enhancing security at all costs may be a tall order; they may, in fact, be much more interested in finding optimal security investments, which may not result in improved security.

To complement our study we are interested in studying properties of a network with varying fractions of expert to naïve users. Further, we want to address the desire of some computer experts to sacrifice individual resources to improve system resilience to attacks, by introducing *benevolent* agents. As discussed above, our analysis thus far evidences the need for such benevolent agents. As a practical example, censorship-resilient networks are run by volunteers; without these benevolent participants, the whole network collapses. This paper shows that there is little hope for strong security if all participants are either naïve, or selfish.

To analyze the impact of the different information conditions we have proposed a new mathematical formalization. We measure the value of complete information as the ratio of the payoff in the complete information environment to the payoff in the incomplete information environment. Our analysis of Figure 3 is a first step in that direction, however, a more formal analysis is deferred to future work.

Finally, a system designer is not only interested in the payoffs of the network participants given different information realities (e.g., due to frequent changes in attack trends). He is also concerned with how well-fortified the organization is against attacks. To that effect we plan to include a more thorough presentation of the parameter conditions that cause attacks to fail due to system-wide protection, and when they succeed (due to coordination failures, passivity, and self-insurance).

References

1. A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC'04)*, pages 21–29, New York, NY, May 2004.
2. A. Acquisti, R. Dingledine, and P. Syverson. On the economics of anonymity. In *Proceedings of the Seventh International Conference on Financial Cryptography and Data Security (FC'03)*, pages 439–443, Gosier, Guadeloupe, Jan. 2003.
3. A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, January–February 2005.
4. A. Acquisti and H. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, Summer 2005.
5. R. Anderson. Liability and computer security: Nine principles. In *Proceedings of the Third European Symposium on Research in Computer Security (ESORICS 1994)*, pages 231–245, Brighton, UK, Nov. 1994.
6. R. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley Computer Publishing, New York, NY, second edition, 2001.
7. T. August and T. Tunca. Network software security and user incentives. *Management Science*, 52(11):1703–1720, Nov. 2006.
8. D. Besnard and B. Arief. Computer security impaired by legitimate users. *Computers & Security*, 23(3):253–264, May 2004.
9. R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of the Fifth Annual Workshop on the Economics of Information Security (WEIS'06)*, Cambridge, UK, June 2006.
10. J. Bolot and M. Lelarge. A new perspective on internet security using insurance. In *Proceedings of the 27th Conference on Computer Communications (INFOCOM'08)*, pages 1948–1956, Phoenix, AZ, Apr. 2008.
11. K. Campbell, L. Gordon, M. Loeb, and L. Zhou. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448, 2003.
12. H. Cavusoglu, S. Raghunathan, and W. Yue. Decision-theoretic and game-theoretic approaches to IT security investment. *Journal of Management Information Systems*, 25(2):281–304, Fall 2008.
13. 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, Aug. 2004.
14. D. Dörner. *The Logic Of Failure: Recognizing And Avoiding Error In Complex Situations*. Metropolitan Books, 1996.
15. A. Etzioni. On thoughtless rationality (rules-of-thumb). *Kyklos*, 40(4):496–514, Nov. 1987.
16. R. Frank. Shrewdly irrational. *Sociological Forum*, 2(1):21–41, Dec. 1987.
17. N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In *Proceedings of the 13th International Conference Financial Cryptography and Data Security (FC'09)*, Christ Church, Barbados, Feb. 2009.
18. E. Gal-Or and A. Ghose. The economic incentives for sharing security information. *Information Systems Research*, 16(2):186–208, June 2005.
19. L. Gordon and M. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, Nov. 2002.
20. L. Gordon, M. Loeb, and W. Lucyshyn. Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy*, 22(6):461–485, Nov. 2003.
21. 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, Apr. 2008.
22. J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*, pages 160–169, Chicago, IL, July 2008.
23. J. Grossklags, B. Johnson, and N. Christin. When information improves information security. Technical report, UC Berkeley & Carnegie Mellon University, CyLab, Feb. 2009. Available at `http://www.cylab.cmu.edu/research/techreports/tr_cylab09004.html`.
24. C. Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of the New Security Paradigms Workshop*, Oxford, UK, Sept. 2009.

25. C. Jaeger, O. Renn, E. Rosa, and T. Webler. *Risk, uncertainty, and rational action*. Earthscan Publications, London, UK, 2001.
26. Kabooza. Global backup survey: About backup habits, risk factors, worries and data loss of home PCs, Jan. 2009. Available at: `http://www.kabooza.com/globalsurvey.html`.
27. D. Kahneman and A. Tversky. *Choices, values and frames*. Cambridge University Press, Cambridge, UK, 2000.
28. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, Mar. 2003.
29. M. Lettau and H. Uhlig. Rules of thumb versus dynamic programming. *American Economic Review*, 89(1):148–174, Mar. 1999.
30. Y. Liu, C. Comaniciu, and H. Man. A bayesian game approach for intrusion detection in wireless ad hoc networks. In *Proceedings of the Workshop on Game Theory for Communications and Networks*, Pisa, Italy, Oct. 2006. Article No. 4.
31. J. McCalley, V. Vittal, and N. Abi-Samra. Overview of risk based security assessment. In *Proceedings of the 1999 IEEE PES Summer Meeting*, pages 173–178, July 1999.
32. D. Meier, Y. Oswald, S. Schmid, and R. Wattenhofer. On the windfall of friendship: Inoculation strategies on social networks. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*, pages 294–301, Chicago, IL, July 2008.
33. A. Miura-Ko, B. Yolken, J. Mitchell, and N. Bambos. Security decision-making among interdependent organizations. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 66–80, Pittsburgh, PA, June 2008.
34. T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC'06)*, pages 35–44, Denver, CO, July 2006.
35. NCSA/Symantec. Home user study, Oct. 2008. Available at: `http://staysafeonline.org/`.
36. P. Paruchuri, J. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *Proceedings of the 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 895–902, Estoril, Portugal, May 2008.
37. I. Png, C. Wang, and Q. Wang. The deterrent and displacement effects of information security enforcement: International evidence. *Journal of Management Information Systems*, 25(2):125–144, Fall 2008.
38. S. Radosavac, J. Kempf, and U. Kozat. Using insurance to increase internet security. In *Proceedings of the 2008 Workshop on the Economics of Networks, Systems, and Computation (NetEcon'08)*, pages 43–48, Seattle, WA, Aug. 2008.
39. S. Schechter and M. Smith. How much security is enough to stop a thief? In *Proceedings of the Seventh International Financial Cryptography Conference (FC'03)*, pages 122–137, Gosier, Guadeloupe, Jan. 2003.
40. J. Stanton, K. Stam, P. Mastrangelo, and J. Jolton. Analysis of end user security behaviors. *Computers & Security*, 2(24):124–133, Mar. 2005.
41. D. Straub. Effective IS Security: An Empirical Study. *Information Systems Research*, 3(1):255–276, Sept. 1990.
42. R. Telang and S. Wattal. An empirical analysis of the impact of software vulnerability announcements on firm stock price. *IEEE Transactions on Software Engineering*, 33(8):544–557, 2007.
43. H. Varian. System reliability and free riding. In L. 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.
44. A. Yannacopoulos, C. Lambrinoudakis, S. Gritzalis, S. Xanthopoulos, and S. Katsikas. Modeling privacy insurance contracts and their utilization in risk management for ICT firms. In *Proceedings of the 13th European Symposium on Research in Computer Security (ESORICS 2008)*, pages 207–222, Málaga, Spain, Oct. 2008.

A Appendix: Tabulated results

Table 1. Weakest link security game: Payoffs for different strategies under different information conditions

Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
$c < b$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$ and $\min_{j \neq i} p_j < b/L$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$M - p_i L$	$M - c$	$M - b$
$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b - p_i L$
$b \leq c$	Incomplete	$M - p_i L$	$M - c$	$M - b - p_i L \left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$

Table 2. Weakest link security game: Conditions to select protection, self-insurance or passivity strategies

Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
$c < b$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c$ and $\min_{j \neq i} p_j < \frac{b}{L}$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c$ and $\frac{b}{L} \leq \min_{j \neq i} p_j$	Complete	$p_i < \frac{b}{L}$	NEVER!	$p_i \geq \frac{b}{L}$
$c < b$	Incomplete	$p_i < \frac{c}{L}$	$p_i > \frac{c}{L}$	NEVER!
$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	Incomplete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c$	Incomplete	$p_i < \frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}}$	$p_i > \frac{c-b}{L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$	$\frac{b}{L\left(1-\frac{b}{L}\right)^{N-1}} \leq p_i \leq \frac{c-b}{L\left(1-(1-\frac{b}{L})^{N-1}\right)}$

Table 3. Weakest link security game: Total expected game payoffs, not conditioned on other players

Case	Information Type	Total Expected Payoff for player i (not conditioned on other players)
$c < b$	Complete	$M - c + \frac{c^2}{2L}$
$b \leq c$	Complete	$M - c + \frac{c^2}{2L} + (c-b)\left(1 - \frac{c+b}{2L}\right)\left(1 - \frac{b}{L}\right)^{N-1}$
$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
$b \leq c \leq \frac{b}{\left(1-\frac{b}{L}\right)^{N-1}}$	Incomplete	$M - c + \frac{c^2}{2L}$
$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < c < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$	Incomplete	$M - c + \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}} + \frac{(c-b)^2}{2L\left(1-\left(1-\frac{b}{L}\right)^{N-1}\right)}$
$\frac{b}{\left(1-\frac{b}{L}\right)^{N-1}} < b + L\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) \leq c$	Incomplete	$M - b - \frac{L}{2}\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right) + \frac{b^2}{2L\left(1-\frac{b}{L}\right)^{N-1}}$
$c < b$	Naive	$M - c + \frac{c^2}{2L}$
$b \leq c$	Naive	$M - b + \frac{b^2}{2L} - \frac{L}{2}\left(1 - \frac{b^2}{L^2}\right)\left(1 - \left(1 - \frac{b}{L}\right)^{N-1}\right)$

Table 4. Best shot security game: Payoffs for different strategies under different information conditions

Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
$c < b$	Complete	$M - p_i L$	$M - c$	$M - b$
$b \leq c$ and $\max_{j \neq i} p_j < \frac{b}{L}$	Complete	$M - p_i L$	$M - c$	$M - b$
$b \leq c$ and $\frac{b}{L} \leq \max_{j \neq i} p_j$	Complete	M	$M - c$	$M - b$
$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b$
$b \leq c$	Incomplete	$M - p_i L\left(\frac{b}{L}\right)^{N-1}$	$M - c$	$M - b$

Table 5. Best shot security game: Conditions to select protection, self-insurance or passivity strategies

Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
$c < b$	Complete	$p_i < c/L$	$p_i \geq c/L$	NEVER!
$b \leq c$ and $\max_{j \neq i} p_j < b/L$	Complete	$p_i < b/L$	NEVER!	$p_i \geq b/L$
$b \leq c$ and $b/L \leq \max_{j \neq i} p_j$	Complete	ALWAYS!	NEVER!	NEVER!
$c < b$	Incomplete	$p_i < c/L$	$p_i \geq c/L$	NEVER!
$b \leq c$	Incomplete	ALWAYS!	NEVER!	NEVER!

Table 6. Best shot security game: Total expected game payoffs, not conditioned on other players

Case	Information Type	Total Expected Payoff
$c < b$	Complete	$M - c + \frac{c^2}{2L}$
$b \leq c$	Complete	$M - b\left(1 - \frac{b}{2L}\right)\left(\frac{b}{L}\right)^{N-1}$
$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
$b \leq c$	Incomplete	$M - \frac{L}{2}\left(\frac{b}{L}\right)^{N-1}$
$c < b$	Naive	$M - c + \frac{c^2}{2L}$
$b \leq c$	Naive	$M - b + \frac{b^2}{2L}$

Table 7. Total effort security game: Payoffs for different strategies under different information conditions

Case	Information Type	Payoff Passivity	Payoff Self-Insurance	Payoff Protection
$c < b$	Complete	$M - p_i L$	$M - c$	$M - b - p_i L\,(1 - 1/N)$
$b \leq c$	Complete	$M - p_i L\,(1 - K/N)$	$M - c$	$M - b - p_i L\,(1 - (K+1)/N)$
$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b - p_i L\,(1 - 1/N)$
$b \leq c$	Incomplete	$M - p_i\,(b + (L-b)/N)$	$M - c$	$M - b - p_i\,(b - b/N)$

Table 8. Total effort security game: Conditions to select protection, self-insurance or passivity strategies

Case	Information Type	Conditions Passivity	Conditions Self-Insurance	Conditions Protection
$c < b$	Complete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c \leq b(N-K)$	Complete	$p_i < \frac{c}{L\left(1 - \frac{K}{N}\right)}$	$p_i \geq \frac{c}{L\left(1 - \frac{K}{N}\right)}$	NEVER!
$b(N-K) < c$	Complete	$p_i < \frac{bN}{L}$	$p_i > \frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}$	$\frac{bN}{L} \leq p_i \leq \frac{c-b}{L\left(1 - \frac{K+1}{N}\right)}$
$c < b$	Incomplete	$p_i < \frac{c}{L}$	$p_i \geq \frac{c}{L}$	NEVER!
$b \leq c \leq b + \frac{b^2}{L}(N-1)$	Incomplete	$p_i < \frac{c}{b + \frac{L-b}{N}}$	$p_i \geq \frac{c}{b + \frac{L-b}{N}}$	NEVER!
$b + \frac{b^2}{L}(N-1) < c$	Incomplete	$p_i < \frac{bN}{L}$	$p_i > \frac{c-b}{b - \frac{b}{N}}$	$\frac{bN}{L} \leq p_i \leq \frac{c-b}{b - \frac{b}{N}}$

Table 9. Total effort security game: Total expected game payoffs, not conditioned on other players

Case	Information Type	Total Expected Payoff
$c < b$	Complete	$M - c + \frac{c^2}{2L}$
$bN \leq L$ and $b \leq c$	Complete	$\sum_{k=0}^{\lfloor N - \frac{c}{b} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1 - \frac{k}{N}\right)} \right)$ $+ \sum_{k=\lfloor N - \frac{c}{b} + 1 \rfloor}^{\lfloor N - 1 - \frac{N}{L}(c-b) \rfloor} Pr[k] \cdot \left(M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2L\left(1 - \frac{k+1}{N}\right)} \right)$ $+ \sum_{k=\lfloor N - \frac{N}{L}(c-b) \rfloor}^{N-1} Pr[k] \cdot \left(M - b - \frac{L}{2}\left(1 - \frac{k+1}{N}\right) + \frac{b^2 N}{2L} \right)$
$L < bN$ and $b \leq c$	Complete	$\sum_{k=0}^{\lfloor N - \frac{cN}{L} \rfloor} Pr[k] \cdot \left(M - c + \frac{c^2}{2L\left(1 - \frac{k}{N}\right)} \right)$ $+ \sum_{k=\lfloor N - \frac{cN}{L} + 1 \rfloor}^{N-1} Pr[k] \cdot \left(M - \frac{L}{2N}(N - k) \right)$
$c < b$	Incomplete	$M - c + \frac{c^2}{2L}$
$bN \leq L$ and $b \leq c \leq b + \frac{b^2}{L}(N-1)$	Incomplete	$M - c + \frac{c^2}{2\left(b + \frac{L-b}{N}\right)}$
$bN \leq L$ and $b + \frac{b^2}{L}(N-1) < c < 2b - \frac{b}{N}$	Incomplete	$M - c + \frac{b^2 N}{2L} + \frac{(c-b)^2}{2\left(b - \frac{b}{N}\right)}$
$bN \leq L$ and $2b - \frac{b}{N} \leq c$	Incomplete	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2 N}{2L}$
$L < bN$ and $b \leq c < b + \frac{L-b}{N}$	Incomplete	$M - c + \frac{c^2}{2\left(b + \frac{L-b}{N}\right)}$
$L < bN$ and $b + \frac{L-b}{N} \leq c$	Incomplete	$M - \frac{1}{2}\left(b + \frac{L-b}{N}\right)$
$c < b$	Naive	$M - c + \frac{c^2}{2L}$
$b \leq c$	Naive	$M - b - \frac{1}{2}\left(b - \frac{b}{N}\right) + \frac{b^2}{L}\left(1 - \frac{1}{2N}\right)$