

When Information Improves Information Security^{*}

(Short paper)

Jens Grossklags^a, Benjamin Johnson^b, and Nicolas Christin^b

^aCenter for Information Technology Policy, Princeton University

^bCyLab, Carnegie Mellon University

jensg@princeton.edu

{johnsonb, nicolasc}@andrew.cmu.edu

Abstract. This paper presents a formal, quantitative evaluation of the impact of bounded-rational security decision-making subject to limited information and externalities. We investigate a mixed economy of an individual rational expert and several naïve near-sighted agents. We further 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.

Key words: Game Theory, Security Economics, Bounded Rationality, Limited Information

1 Introduction

Users frequently fail to deploy, or upgrade security technologies, or to carefully preserve and backup their valuable data [10, 12], 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 [11, 14]. 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 [9].

^{*} We thank John Chuang for his helpful comments to an earlier version of this paper. This work is supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, by the National Science Foundation under ITR awards ANI-0331659 (100x100) and CCF-0424422 (Team for Research in Ubiquitous Secure Technology), and by a University of California MICRO project grant in collaboration with DoCoMo USA Labs.

In prior work, we addressed different canonical types of interdependencies, but we focused our analysis on users capable of gathering all relevant information and correctly understanding all implications of interconnectedness [4, 5]. This paper extends our prior analysis by relaxing several restrictive assumptions on users’ rationality and information availability.

First, 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 [1]. 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.

Second, we 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.

As an initial step, we study the strategic optimization behavior from the perspective of a sophisticated user in an economy of inexperienced users, using three canonical security games that account for externalities [4]. This approach results in a decision-theoretic model [2, 3, 13]. We present the mathematical formulation and analysis methodology in the following section, which is based on our recent work focusing exclusively on the weakest-link externality [6].

2 Model and Analysis

Basic Model: Consider a game in which each of N network users is responsible for choosing security investments for their individual node. Each player begins the game with an initial endowment M , and suffers a maximum loss of L if a security breach occurs. The risk of a security breach is determined by an exogenous probability $p_i \in [0, 1]$, which for this paper we assume to be uniformly distributed. Security risks can be mitigated in two distinct ways. We denote by *protection* those security investment strategies which benefit the public network (such as installing antivirus software or firewalls), and we denote by *self-insurance* those strategies which benefit only the contributing user (such as keeping private data backups) [4]. The cost of full protection investment is denoted b and the cost of full self-insurance is denoted c . Each player chooses a protection investment level $e_i \in [0, 1]$ and a self-insurance investment level $s_i \in [0, 1]$. The manner in which protection strategies jointly affect players in the network is determined by an aggregate protection function $H(e_1, \dots, e_N)$, of which we will consider three types:

weakest link ($H(e_1, \dots, e_n) = \min_j e_j$), *best shot* ($H(e_1, \dots, e_n) = \max_j e_j$), and *total effort* ($H(e_1, \dots, e_N) = \frac{1}{N} \sum_j e_j$). The utility for player i is given by

$$U(i) = M - p_i L(1 - H(e_1, \dots, e_N))(1 - s_i) - be_i - cs_i. \quad (1)$$

Bounded rationality and limited information: We consider distinct approaches to relax assumptions on user rationality and information availability.

First, we distinguish users based on their treatment of network interdependencies. We say a player is *naïve* if she does not take network interdependency into account in her decisions, i.e., she operates under the assumption that $H(e_1, \dots, e_N) = e_i$. Whereas a player is of the type *expert* if she correctly perceives and understands the interdependent nature of the game. That is, she properly considers the role of H into her payoff function.

Second, we distinguish between the level of information users have about others' risks. We say that a player has *complete information* if she knows all the risk factors associated to the other players. In contrast, a player has *incomplete information* if she knows her own risks but not the risks of other players. Clearly, a naïve player does not use the external risk information (because she is not aware of its effects), but an expert player does. To provide a decision framework for an expert with limited information, we assume that the distribution on risk parameters is known to all players. Thus, an expert with limited information can still conduct an expected cost-benefit analysis to make a strategic decision. This setup allows us to study the extent to which the complete information is beneficial to the players and to the network.

Methodology: The overarching aim of our analysis is to understand how expertise and knowledge affect player payoffs. To this end, our study considers a single expert agent in a field of $N - 1$ naïve agents, subject to one of the two information conditions, and the incentives of the utility function defined above.

We begin by determining, for each of the three games, (best shot, weakest link, and total effort), and for each of the two information conditions, (complete and incomplete), the payoff-maximizing strategy for an expert in that game. These strategy conditions involve the parameters b, c, L, M, N, p_i and in the case of complete information p_j for $j \neq i$.

We next compute an average or expected payoff for the expert as a result of playing this strategy. The expected value is taken with respect to the various probabilities p_i and p_j , which we assume are each drawn independently from the uniform distribution on $[0, 1]$. This final expected payoff is a function of parameters b, c, L, M , and N . The required computations often require us to consider selected parameter orientations as separate cases; but there are a small number of separate cases (at most 6) for each game, and the expected payoff functions and case conditions can be recorded neatly in tables.

We compute these payoff functions for an expert with complete information, for an expert with incomplete information, and also for one of the many naïve players. These functions tell the whole story in terms of how a player in this game, with a given level of knowledge and expertise, will fare in a given parameter configuration of this game. Unfortunately, the functions involve five free variables, and it remains to distill the information for further interpretive analysis.

To begin this part of the process, we fix the parameters related to the initial endowment, M , and the total maximum loss, L by setting $M = L = 1$. The assumption $M = L$ says that an agent can lose her entire endowment if a completely unprotected attack occurs, and may be considered as simply an interpretative statement about the initial endowment. The assumption $L = 1$ generates a relatively simple scaling effect – while it does affect the gross payoffs linearly, the assumption does not, for example, affect the configuration of other parameters that yield the minimum or maximum payoff. After incorporating these assumptions, what remains is a payoff function that depends only on the cost of protection, b , the cost of self-insurance, c , and the number of players, N .

The final step is to isolate the parameter conditions that yield interesting and substantive results. It turns out that there is only a narrow range of parameter configuration in which a substantial payoff difference exists between various agent types, and so we focus our attention on those cases.

3 Results

Due to limited space in this version of the paper, we exemplify this methodology by focusing on the case of an expert with limited information in the best shot game. Complete analytical results for all three games, – strategies and expected payoffs, together with all accompanying derivations – may be found in our companion technical report [8] and in our in-depth discussion of the weakest-link externality [6].

3.1 Analytical Results

Consider the factors influencing the decisions of an expert with incomplete information in the best shot game. First, note that because of the linear, monotonous nature of the utility function given in Eqn. (1), only three strategies are potentially utility-maximizing for player i : passivity ($e_i = 0, s_i = 0$), full protection ($e_i = 1, s_i = 0$), and full insurance ($e_i = 0, s_i = 1$). Any other strategy can be shown to result in sub-optimal payoffs [4, 8].

Now, if player i protects, her payoff is $M - b$; if she insures, her payoff is $M - c$, and if she does neither, then her payoff is $M - p_i L (1 - Pr_{-i}^*)$ where Pr_{-i}^* is the probability that one of the $N - 1$ naïve players protects. Because we know the strategy for naïve player j is to protect if and only if $b \leq p_j L$ and $b \leq c$, we may compute Pr_{-i}^* as $\begin{cases} 0 & \text{if } c < b \\ 1 - (\frac{b}{L})^{N-1} & \text{if } b \leq c \end{cases}$. Hence the payoff for the expert when she does nothing is $M - p_i L$ when $c < b$, and $M - p_i L (\frac{b}{L})^{N-1}$ when $b \leq c$. These values can be found in the lower portion of Table 1, which also records initial payoffs as the results of strategy choices under the complete information condition.

The next step is determine what our expert should do given known values for the parameters. Consider the case $b \leq c$. Here the expert with incomplete information would never choose to insure because protection is cheaper for the same result. On the other hand, the assumptions that $b \leq L \leq 1$ imply that the inequality $M - p_i L (\frac{b}{L})^{N-1} \leq$

Table 1. 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 $\max_{j \neq i} p_j < b/L$	Complete	$M - p_i L$	$M - c$	$M - b$
$b \leq c$ and $b/L \leq \max_{j \neq i} p_j$	Complete	M	$M - c$	$M - b$
$c < b$	Incomplete	$M - p_i L$	$M - c$	$M - b$
$b \leq c$	Incomplete	$M - p_i L (b/L)^{N-1}$	$M - c$	$M - b$

$M - b$ is tautological. Hence the expert will also never protect. We find that in the parameter case $b \leq c$, an expert with incomplete information will always choose to be passive. This result can be found in Table 2 under the appropriate parameter case and information condition.

Table 2. 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

Finally, to determine an expected payoff for the expert with incomplete information in the best shot game with $b \leq c$, we compute the probability that she protects (over her draw of p_i) times the expected payoff for protection, plus the probability that she insures times the expected payoff for insuring, plus the probability that she is passive times her expected payoff for passivity. The end result is $M - \frac{L}{2} \left(\frac{b}{L}\right)^{N-1}$. This value is recorded in Table 3, along with expected total payoffs for other parameter cases and player information conditions.

Tables and derivations for the other two canonical games can be found in the companion technical report [8].

3.2 Applications

We highlight two applications from our study. A more thorough discussion can be found in the technical report [8].

Our first result is that the range of security parameters conducive to an environment in which limited information plays a significant role is somewhat restricted, and

Table 3. Best shot security game: Total expected game payoffs

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




Fig. 1. Strategy boundaries in the incomplete information scenario for the expert player as a function of protection and self-insurance costs. (Here the number of players is fixed at $N = 4$, and the initial endowment and the potential loss are fixed at $M = L = 1$.) For the relevant analytic results used to construct these graphs, please refer to the technical report [8].




Fig. 2. Heat plots showing the extent of payoff discrepancy between the expert with complete information and the expert with incomplete information as a function of protection and self-insurance costs. (As in Fig. 1, $N = 4$, and $M = L = 1$.) For the relevant analytic results used to construct these graphs, please refer to the technical report [8].

becomes more restricted as the number of players increases. Consider the three graphs in Figure 1. These graphs show dividing lines between various dominant strategies as function of the two types of investment costs – protection and self-insurance. Comparing these to the graphs in Figure 2, the direct relationship is clear between the parameter conditions that determine certain dominant strategies and the parameter conditions under which additional information has an effect on an expert player’s payoff. From these graphs we can tell that information plays a substantial role in the smaller-sized cases, but in the other cases (i.e., N large), the additional information has little to no effect. Another striking observation (not shown in these figures) is that for every game, the size of the interesting parameter cases shrinks toward zero as we increase the number of players.

Our second result is that, even in the parameter ranges in which information plays a role, that role is limited. In all games, the maximum percent increase in expected payoff as a result of additional information is around 18%. This occurs in the weakest link game with exactly four players, and under fixed costs of protection and self-insurance that are set to produce maximum information impact. So having information about costs and risks of others is not a dominant factor – in other words, using educated guesses to predict these values does not hurt the bottom line too much. In contrast with this observation, the expected payoff of the naïve player is greatly reduced as a result of his imperfect situational perception. So if network users do not understand the interdependent nature of their security threats, their payoff is greatly reduced.

3.3 Value of information

The degree of darkening in the graphs of Figure 2 allows us to visualize the payoff discrepancy as a result of incomplete information, and is something we might refer to as the value of information. It is nontrivial to arrive at a definitive answer for this quantity’s best measure and is in fact the subject of a full related paper [7], but we consider the following ratio definition as a first step towards the goal of reasonably quantifying the value of information in this context.

$$\frac{\text{Expected payoff of an expert agent in the complete information environment}}{\text{Expected payoff of an expert agent in the incomplete information environment}}$$

4 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 [4, 14]. 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.

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.

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 2 is a first step in that direction, however, we defer a more formal analysis to a companion research paper [7].

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 and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, January–February 2005.
2. H. Cavusoglu, S. Raghunathan, and W. Yue. Decision-theoretic and game-theoretic approaches to IT security investment. *J. Mgt. Info. Sys.*, 25(2):281–304, Fall 2008.
3. L. Gordon and M. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4):438–457, Nov. 2002.
4. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proc. WWW'08*, pages 209–218, Beijing, China, Apr. 2008.
5. J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proc. ACM EC'08*, pages 160–169, Chicago, IL, July 2008.
6. J. Grossklags and B. Johnson. Uncertainty in the weakest-link security game. In *Proc. GameNets 2009*, pages 673–682, Istanbul, Turkey, May 2009.
7. J. Grossklags, B. Johnson, and N. Christin. The price of uncertainty in security games. In *Proc. (online) WEIS 2009*, London, UK, June 2009.
8. J. Grossklags, B. Johnson, and N. Christin. When information improves information security. Tech. rep., UC Berkeley & Carnegie Mellon University, CyLab, Feb. 2009. <http://www.cylab.cmu.edu/research/techreports/tr-cylab09004.html>.
9. C. Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proc. NSPW'09*, Oxford, UK, Sept. 2009.
10. Kabooza. Global backup survey: About backup habits, risk factors, worries and data loss of home PCs, Jan. 2009. <http://www.kabooza.com/globalsurvey.html>.
11. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, Mar. 2003.
12. NCSA/Symantec. Home user study, Oct. 2008. Available at: <http://staysafeonline.org/>.
13. S. Schechter and M. Smith. How much security is enough to stop a thief? In *Proc. IFCA FC'03*, pages 122–137, Gosier, Guadeloupe, Jan. 2003.
14. 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.