

Are Security Experts Useful? Bayesian Nash Equilibria for Network Security Games with Limited Information^{*}

Benjamin Johnson^a, Jens Grossklags^b, Nicolas Christin^a, and John Chuang^c

^aCyLab, Carnegie Mellon University

^bCenter for Information Technology Policy, Princeton University

^cSchool of Information, University of California, Berkeley

Abstract. A common assumption in security research is that more individual expertise unambiguously leads to a more secure overall network. We present a game-theoretic model in which this common assumption does not hold. Our findings indicate that expert users can be not only invaluable contributors, but also free-riders, defectors, and narcissistic opportunists. A direct application is that user education needs to highlight the cooperative nature of security, and foster the community sense, in particular, of higher skilled computer users.

As a technical contribution, this paper represents, to our knowledge, the first formal study to quantitatively assess the impact of different degrees of information security expertise on the overall security of a network.

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

1 Introduction

To what extent does information security expertise help make a network more secure?

Common sense seems to dictate that the more security experts participate in the network, the higher the level of overall security should be, since each expert can contribute her own knowledge to improving the security of all parties. However, such a reasoning does not take into account that most modern information networks such as the Internet are distributed over potentially competing entities. In other words, it may not be in everybody's best interest to contribute resources to secure the network, particularly if the benefits of such contributions are unclear.

As an illustration, consider a simple denial-of-service example. The attacker is a customer of Internet Service Provider (ISP) A , while the target is a customer of a different ISP B . A is not directly connected to B , instead traffic going from A to B may have to cross several other administrative boundaries (e.g., ISPs C , D , ...), causing potential congestion at all of these intermediaries. A very desirable solution, from an engineering standpoint, is to filter traffic at ISP A , before it even enters the rest of the

^{*} This research was partially supported by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by the National Science Foundation under ITR award CCF-0424422 (TRUST).

This work is published in Proc. ESORICS'10. The original publication is available at www.springerlink.com, <http://www.springerlink.com/content/978-3-642-15496-6>

network (*ingress filtering*, [9]). Yet, what is the incentive for ISP A to perform ingress filtering? From A 's perspective, ingress filtering means they have to refuse some of their customers' traffic and perhaps deploy intrusive techniques such as deep packet inspection, in order to improve their competitors' security, which may be an economically questionable proposition. Worse even, with modern attack configurations (e.g., botnets [20]), where attack traffic is originating from several different ISPs (A_1, A_2, \dots, A_n) it may be difficult for the source ISPs A_i to distinguish attacks from benign traffic. Thus, ingress filtering could result in unwillingly discarding legitimate traffic, which in turn would certainly lead to loss of reputation, and of customer revenue at the filtering ISPs.

Through the denial-of-service example, we observe that negative externalities play a predominant role in security strategies, even (and especially) when network participants are knowledgeable about security dependencies. Furthermore, in large scale networks such as the Internet, the limited information available about other players renders the decision choices even more complex.

In our previous work [16, 17], we explored the impact of limited information on the strategies chosen by a single expert in a network of naïve players all facing a similar security threat. Naïve players took a myopic view of the network, ignoring all externalities, while the expert player had better threat modeling, and understood externalities arising from the threat faced; this naïve approach is substantiated by empirical surveys of organizational security, e.g., [4]. Addressing the problem relied on a decision-theoretic formulation, where the expert was optimizing her strategy in absence of strategic competition from the naïve players.

In the present work, we address the more general case, where a number k of experts ($1 \leq k \leq N$, with N the size of the network) are making security decisions based both on the security threat they face, and on the behavior of other (expert and naïve) players. The main contribution of this paper is to give formal elements of answer to the question posed in the preamble to this paper, that is, to provide a formal characterization of the impact of the number of competing expert players on the overall security of an information network.

We build upon previously proposed models of network security games [14, 15, 35], which abstract security interactions between network participants by a set of stylized games (weakest link, best shot, total effort). The strategy set for each player is defined by two actions: self-protection, and self-insurance. Self-protection reduces the probability an attack is successful in causing losses to the player, while self-insurance deterministically reduces these losses when a successful attack has happened. Protection is a public good, in that the level of protection chosen by each player impacts (positively or negatively) the overall level of protection all players receive, while self-insurance is a private good, which only benefits the players investing in it.

The remainder of this paper is organized as follows. We review related work in Section 2. In Section 3, we describe our security model, including several simplifications to models presented in our prior work [14]. In Section 4, we explain the methodology of our game-theoretic analysis. Section 5 discusses our numerical results and we conclude in Section 6.

2 Related work

2.1 Limited information

Strategic aspects of different interdependency scenarios have been studied in limited information environments (without self-insurance [18]). In particular, Xu computes Bayesian Nash equilibria for the traditional best shot game when the net benefit of the public good is common knowledge, however, each player's contribution cost is her private information and the general distribution of effort costs are known to all players [38]. Burnett utilizes the weaker link model (in which marginal and declining benefits of contributions above the group minimum exist [8]) to study a model of defense against invasive species. She also assumes common knowledge about the distribution of costs, and limited information about individual effort costs [5]. Somewhat complementary is Manzini and Mariotti's model on negotiations between individual actors and small alliances [24]. In this model, they presume that alliance members are uninformed about the other members' negotiation toughness.

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 [11, 12]. Empirical papers explore the impact of mandated disclosures [6] or publication of software vulnerabilities [33] on the financial market value of corporations. Other contributions to the security field include the computation of Bayesian Nash outcomes for an intrusion detection game [23], security patrol versus robber avoidance scenarios [27], and the preservation of location privacy in mobile networks [10].

2.2 Mixed player populations and bounded rationality

In the economic literature, the distinction between perfect rational and bounded rational preferences has found recognition in several psychologically-inspired studies. For example, analytical models for present-biased preferences (i.e., where individuals assign stronger relative importance to a moment in time as it gets closer) consider both sophisticated and naïve agents. The former foresee and comprehend their self-control problems, whereas the latter lack the strategic understanding of their personal fallacies [26]. Acquisti studies these types of preferences in the context of information security [1]. In a related paper, Acquisti and Varian study how individuals interact with a price-discriminating merchant. In their work, myopic agents do not react strategically to the seller's actions [3]. An additional form of preferences is the resolute type where agent stick to their initial strategy during each subsequent choice [36].

However, common to these studies is typically that they evaluate the decision-making behavior of the different agent types separately, whereas in practice it is likely that many marketplace decisions are subject to the interaction with mixed populations with diverse degrees of sophistication. An exception are studies that consider the choice of an intermediary given the interaction with a probabilistically unknown agent type. For example, in signalling games an agent is endowed by nature with a particular type

(e.g., good or malicious) and the agent can provide evidence of her type to the intermediary via a costly signal [31]. In a similar fashion, Stigler interpreted the concern for privacy as the rightful ownership of knowledge about a person. Individuals want to achieve and maintain a high degree of reputation while keeping undesirable information a secret [32]. Intermediaries such as employers want to learn the individual’s type via signalling or other forms of information revelation.

In agent-based economics, mixed populations have been mostly studied in the tournament context with the goal to determine the most successful agent types in a challenging environment. For example, in the 2001 Trading Auction Competition agents lumped together in small groups were tasked to arrange and automatically purchase travel bundles [37]. Similarly, in the Santa Fe Double Auction tournament, researchers were invited to submit automatic agents to compete on an auction market [29]. In the auction tournament, a majority of agent types did not follow explicit optimization principles, but rather applied heuristic principles [22]. The trading approaches varied in complexity, and the capability to predict and adapt [29]. A more recent example is the lemonade stand tournament in which agents were able to optimize direct sales aspects (such as inventory, quality, and price) and generate indirect payoffs with information trading [25].

In the online safety context, we might wonder whether security markets or large-scale interactions will wipe out any unfamiliar, non-optimizing psychological phenomena (e.g., naïveté) [28]? We doubt that any current trends are able to overcome the rationality obstacles furthered by the complexity of online interactions, and existing information barriers [2]. Therefore, in absence of an effective and sustainable learning process for computer security we have to consider distributional and efficiency consequences of bounded rationality, and the impact on overall system security.

3 Model overview

The security model adopted in this paper builds upon the hybrid public/private goods model that was defined in [14], and extended and refined in [15, 17]. Specifically, we consider a network of N ($N \in \mathbb{N}$) agents that are subject to an exogenous security threat. Each agent i is endowed with an amount of money M_i , and stands to lose L_i if the attack is successful. For now, we hypothesize that attacks occur with a probability p , exogenous to the game.

Agents can reduce the probability an attack is successful, by choosing a self-protection level e_i ($0 \leq e_i \leq 1$). They can also lower the damage in case an attack is in fact successful, with a self-insurance effort s_i ($0 \leq s_i \leq 1$). Self-protection and self-insurance have nominal costs of b_i and c_i , respectively.

With these definitions, the expected payoff to player i is defined as

$$U_i = M_i - pL_i(1 - s_i)(1 - H(e_1, \dots, e_N)) - b_i e_i - c_i s_i, \quad (1)$$

where H is a joint “contribution” function that characterizes the externalities in the security game being played. As in [35], we consider three different security games:

- *Weakest-link*: The level of protection achieved by player i is that contributed by the least protected player in the entire network. That is, $H(e_1, \dots, e_N) = \min_{1 \leq j \leq N} \{e_j\}$.

- *Best shot*: The level of protection achieved by player i is that contributed by the most-protected player in the network. That is, $H(e_1, \dots, e_N) = \max_{1 \leq j \leq N} \{e_j\}$.
- *Total effort*: The level of protection achieved by player i is equal to the average protection level for the N players. That is, $H(e_1, \dots, e_N) = \frac{1}{N} \sum_{j=1}^N e_j$.

All three games have practical applications [14]. Best shot, for instance, is useful in modeling the level of security achieved in a censorship-resilient network like a mix-net [7], where the overall security of the network is guaranteed as long as a single node remains uncompromised. Weakest link is a relatively accurate modeling of perimeter security, where a network, shielded from the rest of the Internet by connecting only through a set of protected machines (routers, firewalls) can be reached by intruders as long as these intruders manage to compromise *one* of these protected machines. Finally, total effort is a good model of parallelized file transfers (e.g., in the case of a peer-to-peer network), where the achieved throughput is roughly equal to the aggregate throughput provided by all peers one is downloading from.

3.1 Simplifications and additional assumptions

In an effort to alleviate notational burden, we adopt several simplifying modifications to the model described by Eqn. (1). The purpose of the following discussion is to justify that these simplifications do not considerably restrict the generality of our results and conclusions.

First, the total endowment M_i is not relevant to the strategies that a player chooses, as it is an a priori known constant. We will dispense with this parameter, studying instead the net changes in the utility U_i (which we now denote by u_i).

Since we are now considering only changes in utilities, it makes sense to normalize the remaining parameters, so that the maximum total loss to player i is set to 1. Henceforth we assume that $b_i, c_i, L_i \in [0, 1]$.

The next change is to incorporate the attack probability p into the loss parameter L_i , treating L_i an *expected* loss, as opposed to a realized loss. This modification is mathematically without loss of generality since the model was simply treating p as a known constant.

Ultimately, Eqn. (1) becomes:

$$u_i = -L_i(1 - s_i)(1 - H(e_1, \dots, e_N)) - be_i - cs_i, \quad (2)$$

with $0 \leq b_i, c_i, L_i, e_i, s_i \leq 1$, and H as defined above for the three games under consideration.

All remaining assumptions about the model are the same as those adopted in [17]. To simplify the analysis we assume that self-insurance and protection costs are homogeneous.¹ To focus on the interesting case in which protection and self-insurance are rationally interesting prospects we assume that $b_i, c_i \leq L_i$. Due to our intent to focus on utility-maximizing equilibria we assume that e_i, s_i are in fact discrete decision variables taking values in $\{0, 1\}$. Each of these assumptions is discussed more thoroughly in [17].

¹ Note that, for the full information, only-expert player case, we explored the case where b_i and c_i are heterogeneous in [15].

Finally, to improve readability of the presentation, we will derive our initial results on protection equilibrium while ignoring caveats resulting from the availability of self-insurance. Expanded results incorporating self-insurance are included in Appendix A.

4 Analysis

4.1 Methodology

To determine how the composition of experts can affect systemwide network protection in a security context, we analyze three distinct N -player security games in which there are k selfish experts ($0 \leq k \leq N$) and $N - k$ naïve players. We assume that the experts understand the dynamics of each game and choose strategies to maximize their own expected payoffs, while the naïve players choose whether to protect based on a simple cost-benefit analysis – comparing protection costs to their individual expected loss in the event of protection failure.

For expert agents, we distinguish between knowledge conditions about other players' expected losses. In a complete information environment, each expert has full knowledge of the expected losses of all players resulting from network protection failure. In an incomplete information environment, each expert knows her own expected loss in the event of protection failure, but only knows the distribution on these expected losses for all players (the uniform distribution).

Finally, to illustrate the drawbacks of selfish equilibria more generally, we compute the social optimum strategy for each game. To facilitate comparisons with the scenario involving selfish experts, we characterize the social optimum from the individual perspective of N cooperative agents.

4.2 Protection Strategies in the Best Shot Security Game

In the best shot security game, the security of the network is determined by the highest protection level of any player (i.e. $H(e_1, \dots, e_N) = \max_j e_j$). The upshot of this game is that if a single player pays for full protection, then the whole network is protected. This scenario gives selfish experts an incentive to shirk protection responsibilities and pass the buck to other players. In the paragraphs below, we consider three different types of player configurations: k selfish experts with incomplete information, k selfish experts with complete information, and N cooperative experts.

Best Shot: k selfish experts with incomplete information In this player configuration, the rationale for expert i goes as follows. If she protects, she pays b ; and if she does not protect, she pays $L_i \cdot FailE_{\neg i}^* \cdot FailN_{\neg i}^*$, where $FailE_{\neg i}^*$ is the probability (over the distribution on the expected losses L_j) that all other experts fail to protect, and $FailN_{\neg i}^*$ is the probability that all naïve players fail to protect. We can easily compute $FailN_{\neg i}^* = b^{N-k}$ (since naïve player j protects if and only if $L_j \geq b$). It follows that expert i should protect if and only if $b \leq L_i \cdot FailE_{\neg i}^* \cdot b^{N-k}$, or equivalently, $L_i \geq \frac{1}{b^{N-k-1} \cdot FailE_{\neg i}^*}$. To solidify the proper strategy, it remains to determine $FailE_{\neg i}^*$. To do this, we assume that the strategy is a symmetric Bayesian Nash equilibrium. In

this event, all experts have the same (pure) strategy of the form: "protect if and only if $L_j \geq \alpha$," (where α depends on b , k , and N). Since α represents the lower protection threshold for all experts, we have $FailE_{i,i}^* = \alpha^{k-1}$, and we can solve for α using $\alpha = \frac{1}{b^{N-k-1} \cdot \alpha^{k-1}}$. The result is $\alpha = b^{\frac{-N+k+1}{k}}$. Thus the equilibrium strategy for expert i is to protect iff $L_i \geq b^{\frac{-N+k+1}{k}}$.

If $k < N$, then this strategy reduces to the simple conclusion that expert i should *never* protect.² The explanation is that if there are any naïve players in the network, the likelihood that any such player fails to protect is at most the cost of protection. Thus the expected loss of expert i in the event of a protection failure is strictly less than the cost of protecting.

If there are N players, then the above strategy simplifies to "protect if and only if $L_i \geq b^{\frac{1}{N}}$." Uniform adoption of this strategy among all experts yields a viable symmetric Bayesian Nash equilibrium.

Best Shot: k selfish experts with complete information In this player configuration, the rationale for expert i goes as follows. If any of the naïve players draws $L_i > b$ they will protect, so none of the experts needs to (or in their own selfish interest ought to) protect. If none of the naïve players make such draws, then the game reduces to one consisting of all expert players. Because the losses are drawn from a uniform distribution, one of the experts will have the highest draw L_j^* . If $L_j^* \geq b$, then this expert should protect. All other experts can refrain from protecting. If all of the L_j are less than b , then the network will remain unprotected. To summarize, the equilibrium strategy is for expert i to protect if and only if $b \leq L_j < L_i$ for all $j \neq i$.

Best Shot: N cooperative experts with complete information We next consider what happens in the best shot game if experts can cooperate to share wealth in an effort to reduce their overall expected costs. First note that if the network is not protected, then the sum of players' losses is $\sum_{j=1}^N L_j$. If $\sum_{j=1}^N L_j \geq b$, then for the purpose of maximizing expected outcomes, it would be advantageous for each expert to contribute toward ensuring protection. One quite reasonable strategy is for player i to pay $\frac{L_i \cdot b}{\sum_{j=1}^N L_j}$, and for the sum paid (which is b) to be used to pay for a single player's protection cost (say the player with the highest L_j). This strategy is fair, symmetric, and results in the lowest possible sum of expected long term costs for all players.

4.3 Protection Strategies in the Weakest Link Security Game

In the weakest link security game, the security of the network is determined by the lowest protection level of any player (i.e. $H(e_1, \dots, e_N) = \min_j e_j$). The upshot of this game is that for the network to be protected, every player must pay for protection. This scenario gives rational players cause to worry whether other players will defect, ultimately resulting in a notable systemwide protection disincentive. Our analysis of player configurations below is structured analogously to the best shot case.

² This idea confirms the result from [17] which considered the case of one expert.

Weakest Link: k selfish experts with incomplete information In this player configuration, the rationale for expert i can be framed as follows. If she does not protect, then she loses L_i , while if she protects, she pays $b + L_i \cdot (1 - ProtE_{\neg i}^* \cdot ProtN_{\neg i}^*)$, where $ProtE_{\neg i}^*$ is the probability that all other experts protect, and $ProtN_{\neg i}^*$ is the probability that all naive players protect. The condition for player i to choose protection can be expressed as $b + L_i \cdot (1 - ProtE_{\neg i}^* \cdot ProtN_{\neg i}^*) \leq L_i$, and this can be simplified (assuming the probabilities in question are non-zero) to the condition: $L_i \geq \frac{b}{ProtE_{\neg i}^* \cdot ProtN_{\neg i}^*}$.

We know that $ProtN_{\neg i}^* = (1 - b)^{N-k}$, because naive player j protects if and only if $L_j \geq b$; and it remains to determine $ProtE_{\neg i}^*$. As we did in the best shot case, we will assume that the strategy for player i is one component of a symmetric Bayesian Nash equilibrium, and that expert j plays an identical strategy of the form "protect if and only if $L_j \geq \gamma$ ", for some γ depending on b , k , and N . Under these conditions, we have $ProtE_{\neg i}^* = (1 - \gamma)^{k-1}$, and so we may solve for γ using $\gamma = \frac{b}{(1-\gamma)^{k-1}(1-b)^{N-k}}$.

For the purpose of numerical analysis we can rewrite the above equation as

$$\gamma(1 - \gamma)^{k-1} = \frac{b}{(1 - b)^{N-k}}. \quad (3)$$

Unfortunately, there is no algebraic solution for γ when $k \geq 5$. Figure 1 plots γ as a function of b for various values of k and N .




Fig. 1. Evolution of γ as defined by Eqn. (3). We plot the evolution of γ as a function of the protection cost b for various network sizes N and various number of expert users k . Recall that γ is an upper bound for expected losses that determines whether an expert in the given configuration will participate in a protection equilibrium. Player i protects if and only if $L_i \geq \gamma$.

Weakest Link: k selfish experts with complete information In this player configuration, if any naive player j draws $L_j < b$ he will not protect, and so systemwide protection will fail, and so no expert will protect. Similarly if expert j draws $L_j < b$ then (absent cooperation opportunities) she will not protect, and again systemwide protection will fail. The only (noncooperative) scenario in which the network is protected,

thus occurs when each and every player draws $L_j \geq b$. In this case, everyone will protect, and the network will be protected. The equilibrium strategy for expert i is thus to protect if and only if every player j draws $L_j \geq b$.

Weakest Link: N cooperative experts with complete information The cost to protect the entire network in the weakest link game is bN . Thus if $\sum_{j=1}^N L_j \geq bN$, it is, over the long term, advantageous for all the players to cooperate and protect the network. A sensible cooperative equilibrium strategy in this instance is for player i to pay $\frac{L_i bN}{\sum_{j=1}^N L_j}$ if and only if $\sum_{j=1}^N L_j \geq bN$; and for the total amount collected (bN) to be divided equally among all players for the purpose of network protection.

4.4 Protection Strategies in the Total Effort Security Game

In the total effort security game, the security of the network is determined by the average protection level of players, (i.e. $H(e_1, \dots, e_N) = \sum_j \frac{1}{N} e_j$). The upshot of this game is that every player receives a partial benefit from his or her contribution to the protection effort. It turns out that this benefit does not depend on the number of other players who choose to protect. Thus a player's decision to contribute to the protection effort can be made without considering the choices of other players. We discuss three player configuration below, following the format of the other two games described above.

Total Effort: k selfish experts with incomplete information In this configuration, player i 's strategy can be framed as follows. If she protects, she pays $b + L_i \cdot (1 - \frac{ExpProt_{-i}^* + 1}{N})$ where $ExpProt_{-i}^*$ is the expected number of players in the network other than i that choose protection. If she does not protect she pays $L_i \cdot (1 - \frac{ExpProt_{-i}^*}{N})$. Thus player i should protect if and only if $b + L_i \cdot (1 - \frac{ExpProt_{-i}^* + 1}{N}) \leq L_i \cdot (1 - \frac{ExpProt_{-i}^*}{N})$. This inequality simplifies to $L_i \geq bN$. Notably, this condition does not depend on the value of $ExpProt_{-i}^*$ or any other variable related to the choices of other players. Expert i should protect if and only if $L_i \geq bN$.

Total Effort: k selfish experts with complete information Similar to the situation above, if experts have complete information in a total effort security game, they can determine other players' protection incentives. However, the core economic structure remains the same as in the case incomplete information. That is, for protection to be a worthwhile investment strategy for player i , it is necessary and sufficient that $L_i \geq bN$.

Total Effort: N cooperative experts with complete information If coordination is allowed in this game, the situation is analogous to the cooperative efforts in the weakest link game. The cost to protect the entire network in the total effort game is still bN . Thus if $\sum_{j=1}^N L_j \geq bN$, it is, over the long term, advantageous for all the players to cooperate and protect the network. A sensible cooperative equilibrium strategy in this instance is for player i to pay $\frac{L_i bN}{\sum_{j=1}^N L_j}$ if and only if $\sum_{j=1}^N L_j \geq bN$; and for the

total amount collected (bN) to be divided equally among all players for the purpose of network protection. Such a strategy, if agreed upon in advance, is guaranteed to yield a higher expected payoff for every player over the course of time – i.e. over the course of numerous draws of expected losses L_i from the uniform distribution on $[0, 1]$.

5 Numerical Illustrations and Observations

To synthesize the information from our previous analysis, we compare, using tables and graphs, the decision outcomes resulting from various configurations of information conditions and expert player configurations. For each game, we compute the conditions for expert i to protect, the probability that expert i protects (over the distribution on L_i), the expected contribution of expert i , and the expected level of network protection – where 1 denotes complete network protection and 0 denotes no protection. All probabilities and expected values are computed assuming each L_i is drawn independently from the uniform distribution on $[0, 1]$. Our discussion of these numerical results focuses primarily on the effect of experts on the probability of network protection.

5.1 Best Shot

Systemwide protection results for the best shot game are shown in Table 1.

Table 1. Best Shot Security Game: Bayesian Nash Symmetric Protection Equilibrium with N Players.

Composition of Expert Players	Experts' Knowledge of Losses	Conditions under which Expert i Protects	Probability that Player i Protects	Expected Contribution from Player i	Probability of Network Protection
k Selfish $1 \leq k < N$	incomplete	Never	0	0	$1 - b^{N-k}$
N Selfish	incomplete	$L_i \geq b^{\frac{1}{N}}$	$1 - b^{\frac{1}{N}}$	$b \cdot \left(1 - b^{\frac{1}{N}}\right)$	$1 - b$
k Selfish $1 \leq k \leq N$	complete	\forall Expert $j \neq i, L_i > L_j$ and \forall Naïve $j, L_j < b$ and $L_i \geq b$	$\frac{b^{N-k}(1-b)}{k}$	$\frac{b^{N+1-k}(1-b)}{k}$	$1 - b^N$
N Naïve	-	$L_i \geq b$	$1 - b$	$b(1 - b)$	$1 - b^N$
N Cooperative	complete	$\sum_i L_i \geq b$	$1 - \frac{b^N}{N!}$	$\frac{b}{N} \left(1 - \frac{b^N}{N!}\right)$	$1 - \frac{b^N}{N!}$

Perhaps the most interesting point to observe about this table in terms of overall network protection is that, in the incomplete information case, increasing the number of experts actually decreases the protection level of the network. This is because experts are incentivized to free-ride when other players are paying for protection. The more experts there are in the network, the more freeriding takes place. This effect can be seen directly in Figure 2 which plots the expected systemwide network protection level as a

function of protection costs for various configurations of experts in a 6-player best shot game.




Fig. 2. Best shot. Evolution of the network protection level as a function of the protection cost b . The different plots vary the number of experts k in a network of $N = 6$ players. We observe that the fewer experts participating in the game, the higher the network protection level is, on average.

In the complete information case, the expected protection outcome for the network is the same regardless of the number of experts. The net effect is always that the network will always be protected if any of the players draws $L_i \geq b$.

The best shot game is especially effective at highlighting the advantage of cooperation. In a configuration of cooperative experts, each player bears a substantially smaller expected cost burden, and the network is far more likely to be secure, in comparison to the analogous configuration of selfish experts with complete information.

5.2 Weakest Link

Systemwide protection results for the weakest link game are shown in Table 2. As was the case in the shot game, the limited information scenario has the property that increasing the number of experts in the game decreases the protection level of the network. Experts are influenced by the risk that other players will be the weak link that causes a protection failure. This risk has a cascading effect, so that as more experts are added, the risk of defection increases.

Except for the similarity between 1 and 2 experts, each additional expert reduces the likelihood of systemwide protection. Figure 3 plots the expected systemwide network protection level as a function of protection costs for various configurations of experts in a 6-player weakest link game.

Observe that in configurations of experts with limited information in the weakest link game, there is an abrupt cut-off in the protection levels facilitating protection conditions. As shown in the decision analysis, once the price of protection exceeds a given value, there no longer exist any protection equilibrium, and so all the experts in the network will choose not to protect.

In the case of complete information, the only scenario in which the network is protected is when, for every player, the individual expected loss is more than the protection

Table 2. Weakest Link Security Game. Bayesian Nash Symmetric Protection Equilibrium with N Players

Composition of Expert Players	Experts' Knowledge of Losses	Conditions under which Expert i Protects	Probability that Player i Protects	Expected Contribution from Player i	Probability of Network Protection
k Selfish $1 \leq k \leq N$	incomplete	* $L_i \geq \gamma$	$1 - \gamma$	$b(1 - \gamma)$	$(1 - b)^{N-k}(1 - \gamma)^k$
k Selfish $1 \leq k \leq N$	complete	$\forall j, L_j \geq b$	$(1 - b)^N$	$b(1 - b)^N$	$(1 - b)^N$
N Naïve	-	$L_i \geq b$	$(1 - b)$	$b(1 - b)$	$(1 - b)^N$
N Cooperative	complete	$\sum_j L_j \geq bN$	** ρ	$b\rho$	ρ

* γ is the least positive solution (if a solution exists) to the equation $\gamma(1 - \gamma)^{k-1} = \frac{b}{(1-b)^{N-k}}$.
 ** ρ can be computed using well-known formulas for the uniform sum distribution. If $bN \leq 1$, we have $\rho = 1 - \frac{(bN)^N}{N!}$. More generally, $\rho = \int_{bN}^N \frac{1}{2(N-1)!} \sum_{k=0}^N (-1)^k \binom{N}{k} (bN - k)^{N-1} \cdot \text{sgn}(bN - k)$.

cost. The same network protection outcome results from a configuration of N naïve players (although the naïve players pay a much higher expected cost for the same net protection effect).


In the cooperative game, the expected protection cost for each player is a bit higher, but the overall expected cost is less (compared to the analogous game with selfish experts), and systemwide network protection is substantially improved. Computing the expected protection contribution for this game requires determining the likelihood that a sum of independently and uniformly distributed random variables from $[0, 1]$ exceeds an arbitrary threshold (bN). The desired probability is easily computed using well-known formulas for the uniform sum distribution, although it is somewhat cumbersome to express.

5.3 Total Effort

Systemwide protection results for the total effort game are shown in Table 3. This game differs from the best shot and weakest link games in that the decision to protect does not depend on the choices of the other players. It is individually worthwhile for an expert to protect the network only if the cost of protection is a $\frac{1}{N}$ fraction of the player's expected loss. For experts, this results in a high threshold for protection – an unfortunate occurrence since protection by any individual player would increase the utility for every player.

In the configuration consisting only of naïve players, protection is much more likely, even though much of the time (over the distribution on L_i) paying that protection cost is a losing investment. This can be seen by comparing the naïve configuration to the cooperative one as shown in Figure 4. The expected network protection level for naïve users exceeds the social optimum protection level whenever $b \geq \frac{1}{2}$.

The cooperative game affords the same result as in the weakest link game.



(a) Evolution of expected network protection as a function of b . The thick lines represent cases on which we zoom in (b). (b) Zoom on small values of the protection cost b .

Fig. 3. Weakest link. Evolution of the network protection level as a function of the protection cost b . The short lines illustrate the presence of limiting conditions on protection equilibria for this game. Where the lines end, the expected network protection level becomes zero. Also note that the cases $k = 1$ and $k = 2$ produce identical curves.

Table 3. Total Effort Security Game. Bayesian Nash Symmetric Protection Equilibrium with N Players

Composition of Expert Players	Experts' Knowledge of Losses	Conditions under which Expert i Protects	Probability that Player i Protects	Expected Contribution from Player i	Expected Level of Network Protection
k Selfish $1 \leq k \leq N$	incomplete	$L_i \geq bN$	$1 - bN$	$b(1 - bN)$	$1 - bN$ if $bN < 1$
k Selfish $1 \leq k \leq N$	complete	$L_i \geq bN$	$1 - bN$	$b(1 - bN)$	$1 - bN$ if $bN < 1$
N Naïve	-	$L_i \geq b$	$(1 - b)$	$b(1 - b)$	$(1 - b)$
N Cooperative	complete	$\sum_j L_j \geq bN$	* ρ	$b\rho$	ρ

* ρ can be computed using well-known formulas for the uniform sum distribution. If $bN \leq 1$, we have $\rho = (1 - \frac{(bN)^N}{N!})$. More generally, $\rho = \int_{bN}^N \frac{1}{2(N-1)!} \sum_{k=0}^N (-1)^k \binom{N}{k} (bN - k)^{N-1} \cdot \text{sgn}(bN - k)$.




Fig. 4. Total effort. Evolution of the network protection level as a function of the protection cost b . For any number of experts $k \geq 1$, the network protection level is inferior to that obtained with a network solely consisting of naïve players. The cooperative equilibrium, here, provides a less desirable overall system outcome as soon as b exceeds 0.5.

6 Discussion and Conclusion

We carried out a game-theoretic analysis of the impact of the number of security experts in a network of competitive players facing a common security threat, under limited information. Our results are somewhat counter-intuitive, in that in all three scenarios considered (best shot, weakest-link and total effort), the addition of selfish experts actually *never* increases the expected level of systemwide network protection. This outcome is rather unexpected, considering our previous result [17], which showed that a lone expert in a network of naïve players stood to make considerable payoff gains compared to naïve players, regardless of the information condition in which the game was played. We thus could have expected that adding more experts would help the network as a whole, but the force of the negative externalities in play *in all scenarios* actually drives the overall security of the network down.

On the other hand, and much less surprisingly, having N *cooperative* experts improves individual expected payoffs, and dramatically increases the expected level of systemwide network protection (relative to N selfish experts).

In sum, we showed in [17] that user expertise could lead to strong individual benefits even in the case of limited information. The present paper shows that, from a community standpoint, expert users may not only be invaluable contributors, but also free-riders preying on the weaknesses of naïve users. As a result, even networks populated with a large share of sophisticated users will frequently not achieve superior security. One of the direct outcomes of this work is that user education needs to highlight the cooperative nature of security, and heighten the community sense of better educated computer users.

While the model proposed has some limitations, and oversimplifications, we believe that it raises a number of points that warrant further consideration by security practitioners.

6.1 Caveats

Our model of limited information is decidedly simple. Only one parameter is unknown to player i , the expected loss L_j for players $j \neq i$. Even so, we assume that player i knows the probability distribution of L_j .³ In practice, even that probability distribution may be unknown to most players, and the nominal costs of protecting or self-insuring (b and c respectively) may also be heterogeneous, and unknown. We note that, if all external costs and potential losses are unknown, all players may have to resort to naïve strategies since they have no way of estimating the externalities in play. More generally, our crude, binary, distinction between naïve and expert players is much more nuanced in practice, where we likely would find a near-continuum of expertise levels, ranging from the abysmally clueless to the highly cognizant. Nonetheless, we believe that our model captures the salient distinction between users who do understand externalities and the role they play, and those who do not.

6.2 Applications

There are three immediate applications of the mathematical results we obtained.

Tragedy of the commons in best-shot games. Best shot security environments are in theory extremely resilient to security threats, since only one agent needs to contribute vigorously (or act securely) to save the whole network. However, our results indicate that there may be a “tragedy of the commons” situation in play, where none of the agents actually has an interest in fulfilling that role. The interesting point is that limited information exacerbates this phenomenon, something we first identified in prior work [14].

Fragility of weakest-link games. Likewise, we discovered that weakest-link games have an increased fragility in presence of expert players and limited information. Weakest-link games are known to offer a vexing set of negative externalities [14, 15, 35], and they rarely converge to a satisfying equilibrium [34]. Unfortunately, one outcome of this work is that limited information degrades even more the likelihood of reaching an acceptable state when players understand the externalities at hand.

Developing side-payment mechanisms. The main take-away from the mathematical formalism proposed in this paper is that, while selfish experts lead to extremely undesirable protection equilibria in the presence of limited information, cooperative experts can completely change the outcome of a game, *and* cooperation can likely be enforced by simply considering side-payments between users [21]. Determining how these side-payments could be carried out highly depends on the context in which the game is played, but it is not inconceivable to imagine bilateral contractual relationships between pairs of players. For example, between ISPs, such side-payments could be made at the time transit agreements are established.

³ For the sake of simplicity of the presentation, we assumed the uniform distribution, but similar derivations could be carried out for any known, “well-behaved” probability distribution.

6.3 Public policy impact

Beyond the applications outlined above, we believe that our work has a clear public policy impact, on at least two levels.

Inter-agency collaboration. In the United States, and in many other countries, national security is not within the purview of a single governmental agency. While, in the U.S., the Department of Homeland Security was created after 9/11 as a vehicle to centralize national security decisions, the reality is that a multitude of government agencies (CIA, NSA, Pentagon, Coast Guard, Bureau of Alcohol, Tobacco and Firearms etc.) are involved, at one degree or another, in such decisions. Such agencies are usually competing for funding budgets, and could be viewed as competing players in national security matters. What we have seen from our model, is that having qualified personnel in all of these agencies actually exacerbates the harmful effects of competition. As such, enforcing collaboration between the different agencies playing a role in containing a threat is not just a desirable goal, it is an absolute necessity to achieve any level of security.

User education The need for technical user education and training has been repeatedly emphasized by security researchers, based on empirical field studies (e.g., [30]) and analytical results (e.g., [16, 17]). Yet, while user education has primarily focused on better understanding of the threats faced, our results indicate that, an equally, if not more important aspect of user education should be in highlighting the cooperative nature of security, and fostering the community sense of high-skilled computer users.

6.4 Future research directions

This paper closes an arc of research on game-theoretic modeling of negative externalities in information security, by complementing our previous work on game-theoretic models of complete information [14, 15], and decision-theoretic models of incomplete information [16, 17]. This does not mean that there are no other interesting directions to pursue on these models, but we believe that to go to the next level, we need to inform our models with field data. In particular, we are interested in removing assumptions on the nature of the various cost functions on which we rely. For instance, surely, the cost incurred by deploying protection measures is not perfectly linear in the degree of protection deployed. Behavioral biases (e.g., risk-aversion [19]) can also create a dichotomy between actual and perceived costs. We have started preliminary investigations through user studies [13], but plan on enriching our models with field data applied to specific contexts. For instance, we are considering whether we can find evidence of correlation between (de)peering decisions and the amount of undesirable traffic (e.g., attack traffic) traversing different ISP networks.

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 and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 3(1):26–33, January–February 2005.
3. A. Acquisti and H. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, Summer 2005.
4. M. Bashir and N. Christin. Three case studies in quantitative information risk analysis. In *Proceedings of the CERT/SEI Making the Business Case for Software Assurance Workshop*, pages 77–86, Pittsburgh, PA, September 2008.
5. K. Burnett. Introductions of invasive species: Failure of the weaker link. *Agricultural and Resource Economics Review*, 35(1):21–28, April 2006.
6. 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.
7. D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, February 1981.
8. R. Cornes. Dyke maintenance and other stories: Some neglected types of public goods. *Quarterly Journal of Economics*, 108(1):259–271, February 1993.
9. P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing, January 1998. RFC 2267.
10. J. Freudiger, M. Manshaei, J.-P. Hubaux, and D. Parkes. On non-cooperative location privacy: A game-theoretic analysis. In *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS'09)*, pages 324–337, Chicago, IL, November 2009.
11. E. Gal-Or and A. Ghose. The economic incentives for sharing security information. *Information Systems Research*, 16(2):186–208, June 2005.
12. L.A. 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, November 2003.
13. J. Grossklags, N. Christin, and J. Chuang. Predicted and observed behavior in the weakest-link security game. In *Proceedings of the 2008 USENIX Workshop on Usability, Privacy and Security (UPSEC'08)*, San Francisco, CA, April 2008.
14. 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.
15. 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.
16. J. Grossklags, B. Johnson, and N. Christin. The price of uncertainty in security games. In *Proceedings (online) of the Eighth Workshop on the Economics of Information Security (WEIS)*, London, UK, June 2009.
17. J. Grossklags, B. Johnson, and N. Christin. When information improves information security. In *Proceedings of the 2010 Financial Cryptography Conference (FC'10)*, Canary Islands, Spain, January 2010.
18. J. Hirshleifer. From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, 41(3):371–386, January 1983.
19. D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, XLVII:263–291, 1979.
20. 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.
21. M. Katz and C. Shapiro. Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440, June 1985.

22. M. Lettau and H. Uhlig. Rules of thumb versus dynamic programming. *American Economic Review*, 89(1):148–174, March 1999.
23. 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*, page Article No. 4, 2006.
24. P. Manzini and M. Mariotti. Alliances and negotiations: An incomplete information example. *Review of Economic Design*, 13(3):195–203, September 2009.
25. A. Noy, D. Raban, and G. Ravid. Testing social theories in computer-mediated communication through gaming and simulation. *Simulation & Gaming*, 37(2):174–194, June 2006.
26. T. O’Donoghue and M. Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, March 1999.
27. 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 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 895–902, Estoril, Portugal, May 2008.
28. M. Rabin. A perspective on psychology and economics. *European Economic Review*, 46(4–5):657–685, May 2002.
29. J. Rust, J. Miller, and R. Palmer. Characterizing effective trading strategies: Insights from a computerized double auction tournament. *Journal of Economic Dynamics and Control*, 18(1):61–96, January 1994.
30. S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The design and evaluation of a game that teaches people not to fall for Phish. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS’07)*, pages 88–99, Pittsburgh, PA, 2007.
31. A. Spence. Job market signaling. *Quarterly Journal of Economics*, 3(87):355–374, August 1973.
32. G. Stigler. An Introduction to Privacy in Economics and Politics. *The Journal of Legal Studies*, 4(9):623–644, December 1980.
33. 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.
34. J. Van Huyck, R. Battalio, and R. Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–248, March 1990.
35. 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.
36. L. von Auer. Revealed preferences in intertemporal decision making. *Theory and Decision*, 56(3):269–290, May 2004.
37. M. Wellman, P. Wurman, K. O’Malley, R. Banger, S. Lin, D. Reeves, and W. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5(2):43–51, March 2001.
38. X. Xu. Group size and the private supply of a best-shot public good. *European Journal of Political Economy*, 17(4):897–904, November 2001.

A Self-Insurance Considerations

In this section, we briefly revisit the prior analysis by considering the ways in which self-insurance further decreases the protection likelihood for expert players. Because some of the derivations required in the complete analysis are especially cumbersome,

especially in the weakest link game, we shall resort to a high level overview of the situation.

A.1 Self-Insurance in the Best Shot Game

In the best shot game, self-insurance is easy to address. Self-insurance is only a spoiler for a self-protection equilibrium when it is cheaper than protection ($c < b$). In this event, all selfish expert players (and even the naïve players) would defect to the self-insurance strategy to improve their own payoff. On the other hand, cooperative experts could still work together to protect the network as long as $b < cN$ (and $\sum_{j=1}^N L_j \geq b$). If $b < c$, then no player will choose the self-insurance strategy because it is cheaper to protect for the same individual result.

A.2 Self-Insurance in the Weakest Link Game

How could the existence of self-insurance spoil the weakest-link protection equilibrium? For a short but not entirely simple answer, the expert i will defect to the self-insurance strategy if $L_i \geq \frac{c-b}{1-Prot_{-i}^*}$ where $Prot_{-i}^*$ is the probability that (under the current value of cost parameters) all the players other than i will protect. Unfortunately, unless the number of other experts is zero, the value of $Prot_{-i}^*$ is not amenable toward a closed form formula involving b , c , N and k . Because self-insurance introduces an *upper bound* on expected losses, a symmetric equilibrium strategy in which the network has a chance of protection requires expert i to protect iff $\alpha \leq L_i \leq \beta$, for some parameters α and β which both depend on b , c , k , and N . Even when $N = k$, determining α and β requires solving the following parametrized system of equations for α and β :

$$\frac{b}{L(\beta - \alpha)^{N-1}} = \alpha \quad (4)$$

$$\frac{c - b}{L(1 - (\beta - \alpha)^{N-1})} = \beta. \quad (5)$$

We can do this, and derive some relations between parameters, but in the end of this process the resulting inequality conditions do not yield substantial insights beyond what is already obvious from a high-level view – namely that availability of self-insurance can be a serious distractor for a protection equilibrium outcome in the weakest link game. We already deduced from the restricted game without self-insurance, that for any protection equilibria to exist, the protection costs must be very small, on the order of a constant times $\frac{1}{N}$. The presence of self-insurance exacerbates the problem, ruining the chances of there being a protection equilibrium even for low values of b and moderately high values of c . (For a simple example, even if b is as low as $\frac{1}{eN}$, self-insurance can serve as a deterrent to protection investment with values of c as high as $\frac{2}{3}$).

A.3 Self-Insurance in the Total Effort Game

In the total effort game, conditions under which the presence of the self-insurance alternative spoils a protection equilibrium are simple to express (although the derivation

of these conditions is nontrivial). If $c < bN$, then there is no protection equilibrium, because an expert for whom it is worthwhile to protect must have a loss L_i exceeding bN . Under such conditions, one can derive that it is more advantageous for the expert to pay the self-insurance premium c . On the other hand, if $c > bN$, then the equilibrium strategy defined previously in which expert i protects if and only if $L_i \geq bN$ continues to be a Pareto-dominant Bayesian Nash equilibrium.

A.4 Overall Effects of Self-Insurance

The results in this section indicate that the existence of self-insurance strategies can be a deterrent toward protection investments, especially in the weakest link game. Indeed the existence of self-insurance only contributes further to the overarching theme of our analysis – that the number of experts does not improve the security of the network when those agents are acting in their own best interest.