Blue versus Red: Towards a Model of Distributed Security Attacks

Neal Fultz* and Jens Grossklags

School of Information, University of California, Berkeley
102 South Hall, 94720 Berkeley, CA
{nfultz,jensg}@ischool.berkeley.edu

Abstract. We develop a two-sided multiplayer model of security in which attackers aim to deny service and defenders strategize to secure their assets. Attackers benefit from the successful compromise of target systems, however, may suffer penalties for increased attack activities. Defenders weigh the force of an attack against the cost of security. We consider security decision-making in tightly and loosely coupled networks and allow defense expenditures in protection and self-insurance technologies.

Keywords: Game Theory, Economics of Security, Tightly and loosely coupled networks, Protection, Self-insurance.

1 Introduction

If you encounter an aggressive lion, stare him down. But not a leopard; avoid his gaze at all costs. In both cases, back away slowly; don't run (Bruce Schneier, 2007 [37]).

The focus of this paper is a better understanding of attacker motives and strategies when faced with diverse defense patterns (i.e., different protection interdependencies). In particular, we want to provide a mathematical framework with enough nuanced structure to enable more intuitive statements about characteristics of cyber-attack equilibria [14]. We add to the literature on game-theoretic models that have often exclusively focused on the strategic aspects of offensive [15,36] or defensive [22,26,30] actions, respectively.[1]

[*] We thank Alvaro Cárdenas, Nicolas Christin, John Chuang, Roger Dingledine, Paul Laskowski, Adam Shostack, Doug Tygar and the anonymous reviewers for their 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.

[1] Several research papers explore the optimal strategies of defenders and attackers in graph-theoretic network inoculation games [4,31]. We explore economic security incentives in different models capturing public goods characteristics and the trade-off between protection and self-insurance.

R. Dingledine and P. Golle (Eds.): FC 2009, LNCS 5628, pp. 167–183, 2009.

The prevalence of widely spread, propagated and correlated threats such as distributed denial of service attacks (DDoS), worms and spam has brought attention to interdependencies existing in computer networks. For an attacker this might create strong economies but sometimes also diseconomies of scale. For example, a single breach of a corporate perimeter may allow an attacker to harvest resources from all machines located within its borders. In other scenarios an attacker may have to shut down every single computer or network connection to achieve an attack goal and thereby incur large costs potentially proportional to network size. More generally, there is an interaction between the structure of the defenders' network, the attack goal and threat model. In Grossklags *et al.* [22] we analyze a set of canonical games that capture some of these interdependencies.

We distinguish between tightly and loosely coupled networks [33]. In a tightly coupled network all defenders will face a loss if the condition of a security breach is fulfilled. This may be a suitable description, for example, of a network perimeter breach that causes the spread of malicious code to all machines, but also applies to independently acting defenders that try to preserve a common secret or resist censorship. In a loosely coupled network consequences may differ for network participants. For example, an attacker might be interested to gain control over a limited set of compromised machines ("zombies" or "bots") and to organize them into a logical network ("botnet") with the goal of executing a DDoS attack against third parties [28] or sending unsolicited information to and from the bots (i.e., popup advertisements and spam). At other times, an attacker might target a specific set of users (e.g., wealthy users in spearphishing scams). Other users would stay unharmed and are never targeted.

With our work we hope to provide a more complete framework to understand defenders' and attackers' incentives and expected security actions and outcomes for a variety of decision-making situations. In the current paper, we are able to discuss which defense actions are plausible given a motivated and strategically acting attacker. We can comment on several important facets of computer security warfare, such as when deterrence will be successful, or when defenders prefer to abstain from any protective action. With our modeling work we expect to provide the foundations for experimental and empirical research, but we are also interested to evolve the model so that it captures more facets of fully distributed attacks.

The rest of the paper is organized as follows. We briefly review related work on models involving strategic attackers and defenders in Section 2. In Section 3 we introduce our game-theoretic model and its relationship to our prior work. We present our analysis in Section 4 and conclude in Section 5.

2 Related Work

In our prior work, we have provided a broader overview of the literature on security economics [22,23]. Our current interest is centered on the incentives of attackers and game-theoretic models with strategically acting defenders and malefactors.

A number of papers provide practical discussions of economic factors related to computer security. Anderson highlights the oftentimes mismatched security incentives between consumers and commercial institutions that host sensitive data or mediate transactions [3]. Franklin *et al.* collect and analyze activity and pricing data from underground marketplaces [18]. Kshetri [29] and Chung *et al.* [12] explore international aspects of cybercrime. Some researchers have conducted survey or interview studies with hackers and cyber-criminals providing rare insights about their motivations and incentives [10,19].

More formally, Schechter and Smith [36] draw upon the economics of crime literature to construct a model of attackers in the computer security context [5]. They derive the penalties and probabilities of enforcement that will deter an attacker who acts as an utility optimizer evaluating the risks and rewards of committing an offense [8]. Similarly, we consider an attacker utility function that allows offensive players to select the force of attacks while they consider potential penalties from enforcement.

Cavusoglu *et al.* [9] analyze 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 responds strategically then in most considered cases the firm will be able to select a more adequate response leading to higher profits. In contrast to Cavusoglu *et al.*, we consider different types of interdependencies and games with multiple attackers and defenders.

Clark and Konrad present a game-theoretic model with one defender and one attacker. The defending player has to successfully protect multiple nodes while the attacker must merely compromise a single point [13]. Their model captures the incentives of a weakest-link game [25], however, with a strategic attacker. We consider multiple individually-rational defenders and allow them to also invest in self-insurance adding an additional perspective to this scenario. Similarly, following Varian's exposition, who also considers strategic attackers, we analyze three canonical contribution functions that determine a common protection level for all defenders [41]. We expand on his analysis of the attacker-defender interaction by considering self-insurance investments as well as security incentives in loosely coupled games.

3 Model

In previous work, we analyzed protection and self-insurance incentives for defenders facing an exogenous attacker [22]. We improve on our *security games* framework by modeling attackers as active and strategic economic actors. In the following, we present the basic framework for the case of N defenders and one attacker. We extend our model to the case of M attackers in Section 4.3.

3.1 Red: Attacker Incentives

The attacker has two actions at her disposal. First, she may choose whether to engage in any attacks at all, and how many defenders k she targets ($0 \leq k \leq N$).

Second, the attacker may choose the force of attacks, a $(0 \leq a \leq 1)$, with $a = 1$ representing the attack with the highest impact. In contrast, $a = 0$ denotes an entirely ineffective and harmless attack strategy. The attacker will receive a benefit that is proportional to the force of her attacks, aL, for each not sufficiently protected defender she is able to compromise.

The attacker has to consider H_e, the group security contribution function of the defenders, which has the decisive impact on whether a targeted defender will be compromised. If $H_e = 1$ the defense efforts will always thwart an attack irrespective of a. A value of $H_e = 0$ leaves the defenders completely vulnerable. We present five different variations of H_e in the section on defender incentives.

Additionally, there is a chance that the attacker is caught and fined F, $F > 0$. The probability of being caught for *each* attack made, p_c, is independent of whether the attack was successful or not. Therefore, the expected utility of attacker i is:

$$Red = \begin{cases} \sum_1^k aL(1 - H_e) - (1 - (1 - p_c)^k)F & \text{if Red attacks } (k > 0), \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

In the current model, we assume that the likelihood of being penalized is related to the number of targeted defenders, k, however, independent of the force of the attack, a. In practice, this likelihood may depend on both parameters since defenders will more frequently involve law enforcement or react vigilantly if attacks are more severe. However, end users and members of small networks are often powerless in their attempts to punish perpetrators of cybercrime. On the one hand, limited and sometimes immaterial damages are an obstacle when users attempt to encourage law enforcement to follow up on their complaints [20]. On the other hand, the cost of identifying an attacker and enforcing a penalty is usually well-beyond the effort needed for a reasonable defense (e.g., cost of forensics, honeypots, maintenance of law enforcement contacts). Users may not want to incur these significant expenses (and we do not consider them in our model). In effect, we assume that a more engaged attacker will face, at least in the aggregate, a higher likelihood of being caught. Of course, there are also obstacles when trying to approximate overall attack activity. For example, enforcement is negatively impacted if multiple jurisdictions are involved [38]. Taken together, we argue that our formulation is a reasonable description for home users and small entities. In contrast, large companies are more likely to mandate thorough investigations and seek involvement of enforcement units after security breaches as a part of their overall security strategy. We defer the analysis of different alternatives for the attacker utility to future work.

3.2 Blue: Defender Incentives

Each of $N \in \mathbb{N}$ defenders receives an endowment W. If she is attacked and compromised successfully, she faces a loss L that is impacted by the force of the attack, a.[2] Defensive players have two security actions at their disposition. Player i can select between a private self-insurance investment, $0 \leq s_i \leq 1$, and a protection level, $0 \leq e_i \leq 1$, that will contribute to a common protection effort. For example, self-insurance includes expenditures in backup technologies, whereas firewalls, patching, and intrusion detections systems are protective efforts [22].[3] Finally, $b \geq 0$ and $c \geq 0$ denote the unit cost of protection and self-insurance, respectively. The generic utility function for a *targeted* defender has the following structure:

$$Blue_i = E(U_i) = W - aL(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i \tag{2}$$

where following usual game-theoretic notation, e_{-i} denotes the set of protection levels chosen by players other than i. If the defender is not targeted (for example, if $k = 0$) then the defender will only incur the cost of protection and self-insurance:

$$Blue_i = E(U_i) = W - be_i - cs_i \tag{3}$$

$H_e = H(e_i, e_{-i})$ is the group "security contribution" function that characterizes the effect of e_i on U_i, subject to the protection levels chosen (contributed) by *all* other players.[4] We will discuss five variations of H_e in the next section.

From Eqs. (2 and 3), the magnitude of a loss depends on three factors: i) whether the defender was targeted by the attacker and with what force of attack (a), ii) whether the individual invested in self-insurance (s_i), and iii) the magnitude of the joint protection level (H_e). Self-insurance always lowers the loss that an individual incurs when compromised by an attack. Protection probabilistically determines whether an attack is successful. Eqn. (2) yields an expected utility.

3.3 Canonical Security Contribution Functions

In prior work [22], we analyzed security games with five different canonical security contribution functions that we will briefly describe in the following. The first three specifications for H_e represent important baseline cases recognized in the public goods literature: total effort, weakest-link and best shot. The attack consequences in these games are tightly coupled; that is, all defenders will face a loss if the level of the security contribution function is not sufficient to block an attack. With two variations of the weakest target contribution function we

[2] For simplicity, we analyze the case where attacker gain and defender loss are identical (if the defender is not self-insured). In practice, we would frequently expect that there is a disparity between the two *subjective* values [2].

[3] We also complement work on market insurance for security and privacy. Cyberinsurance can fulfill several critical functions. For example, audit requirements for cyberinsurance can motivate investments in security, and might contribute to a better understanding of the economic value of the protected resources [27]. Several researchers have investigated the impact of correlation of risks and interdependency of agents in networks on the viability of insurance [6,7,35].

[4] We require that H_e be defined for all values over $(0, 1)^N$. However, we do not place, for now, any further restrictions on the contribution function (e.g., continuity).

analyze security scenarios with loosely coupled attack outcomes. In a loosely coupled network consequences may differ for network participants.[5]

Total/average effort security game (*tightly coupled*): The global protection level of the network depends on the sum of contributions normalized over the number of all participants. That is, we define $H(e_i, e_{-i}) = \frac{1}{N}\sum_i e_i$, so that Eqn. (2) becomes

$$E(U_i) = W - aL(1 - s_i)(1 - \frac{1}{N}\sum_k e_k) - be_i - cs_i \ . \tag{4}$$

With the total effort game we consider, for example, the scenario where an attacker wants to slow down distributed transfer of a file on a P2P network. With fewer users protecting their network connectivity the total efficiency of the data communication will be reduced.

Weakest-link security game (*tightly coupled*): The overall protection level depends on the minimum contribution offered over all entities. That is, we have $H(e_i, e_{-i}) = \min(e_i, e_{-i})$, and Eqn. (2) takes the form:

$$E(U_i) = W - aL(1 - s_i)(1 - \min(e_i, e_{-i})) - be_i - cs_i \ . \tag{5}$$

In the weakest-link scenario an attacker wants to breach the perimeter of a closed network (e.g., a virtual private network) by locating a hidden vulnerability such as a weak password. Similarly, the perpetrator might want to learn the identities of members of a filesharing darknet, or some other secret that is shared between multiple users [14].

Best shot security game (*tightly coupled*): In this game, the overall protection level depends on the maximum contribution offered over all entities. Hence, we have $H(e_i, e_{-i}) = \max(e_i, e_{-i})$, so that Eqn. (2) becomes

$$E(U_i) = W - aL(1 - s_i)(1 - \max(e_i, e_{-i})) - be_i - cs_i \ . \tag{6}$$

Sometimes attackers want to remove from circulation or censor a particular piece of information. In this case, they are participating in a best shot scenario. As long as a single copy remains available to the public domain the attack goal is not achieved [17].

***k*-Weakest-target security game without mitigation (*loosely coupled*):** Here, an attacker will *always* be able to compromise the entities with the k lowest protection levels, but will leave other entities unharmed. This game derives from the security game presented in [11]. Formally, we can describe the game as follows:

$$H(e_i, e_{-i}) = \begin{cases} 0 \text{ if } e_i \leq e_{(k)} \\ 1 \text{ otherwise,} \end{cases} \tag{7}$$

[5] Please refer to our relevant prior work for detailed interpretations of all sub-games [22,23]. Varian [41] and Hirshleifer [25] discuss also applications outside of the security context such as maintenance of dikes on an island.

which leads to

$$E(U_i) = \begin{cases} W - aL(1-s_i) - be_i - cs_i & \text{if } e_i \leq e_{(k)}, \\ W - be_i - cs_i & \text{otherwise.} \end{cases} \tag{8}$$

An attacker might be interested in such a strategy if the return on attack effort is relatively low, e.g., when distributing spam. It is also relevant to an attacker with limited skills, a case getting more and more frequent with the availability of automated attack toolboxes [39]; or, when the attacker's goal is to commandeer the largest number of machines using the smallest investment possible [18].

***k*-Weakest-target security game with mitigation (*loosely coupled*):** This game is a variation on the above weakest target game. Whether an attack on the weakest protected players is successful is now dependent on each target's security level. Here, an attacker is not necessarily assured of success. In fact, if all individuals invest in full protection, not a single machine will be compromised. H_e is defined as:

$$H(e_i, e_{-i}) = \begin{cases} 1 - e_i & \text{if } e_i \leq e_{(k)} \\ 1 & \text{otherwise,} \end{cases} \tag{9}$$

so that

$$E(U_i) = \begin{cases} W - aL(1-s_i)(1-e_i) - be_i - cs_i & \text{if } e_i \leq e_{(k)}, \\ W - be_i - cs_i & \text{otherwise.} \end{cases} \tag{10}$$

This variation of the weakest target contribution function allows us to capture scenarios where, for instance, an attacker targets a specific vulnerability, for which an easily deployable countermeasure exists.

4 Nash Equilibrium Analysis

4.1 One Attacker, One Defender

Let us consider a general defender security function $H_e = H(e)$. For $N = 1, M = 1$, the utility functions are:

$$Blue = \begin{cases} W - aL(1-H_e)(1-s) - be - cs & \text{if Red attacks } (k=1), \\ W - be - cs & \text{otherwise.} \end{cases} \tag{11}$$

$$Red = \begin{cases} aL(1-H_e) - p_cF & \text{if Red attacks } (k=1), \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

We observe that if $p_cF > L$ then the attacker has no incentive to be active $(a = 0, k = 0)$ regardless of the defender's protection decision. On the other hand, if the expected loot (which is subject to the defender's strategy) is greater than the expected fine, a full attack $(a = 1, k = 1)$ dominates other offensive strategies.

If protection is more expensive than self-insurance ($b \geq c$) then the defender has no incentive to protect. Then, self-insurance will be purchased as long as the associated cost is lower than the expected loss ($e = 0$, $s = 1$, if $L > c$ given that $a = 1$).

For an arbitrary contribution function (and $b < c$), interior equilibria may exist and are of the form:

$$H_e = 1 - \frac{p_c F}{L} \tag{13}$$

$$a = \frac{b}{L} \tag{14}$$

These conditions represent an interior solution ($0 \leq (H_e, a) \leq 1; k = 1$) as long as the expected fine for the attacker is not larger than the cost of protection ($p_c F \leq b$), and the loss from a security compromise is at least as large as protection costs ($L \geq b$).

If *only* the first condition delivers a non-permissible value (i.e., $p_c F > L$, but $L \geq b$) then there does not exist a *pure* strategy so that the attacker prefers to be active. That is, when choosing a low attack strength she would evade protection efforts by the defender, however, could not gain enough from the attack to pay for the expected fine. A highly virulent attack would immediately motivate the defender to fully protect. We defer the analysis of mixed strategies for this case to future work.

When the second condition (Eq. 14) does not bind ($L < b$), whether or not Eq. 13 holds, then the defender will remain passive ($e = 0$ and $s = 0$) and he will enable the attacker to successfully compromise his resources ($k = 1$ and $a = 1$ if $p_c F \leq L$).

For the simple contribution function, $H_e = e$, there is no interior solution. However, depending on parameter values there are three simple Nash equilibria: *Passivity*, where the defender does not protect and is attacked; *Full self-insurance* where the defender is attacked but is self-insured; and *Deterrence*, where the attacker does not attack and the defender does not protect.

Result 1: If an interior solution exists, the cost-benefit ratio, b/L, imposes limits on Red's willingness to attack. Therefore, reducing b would lead to less intense attacks and a higher expected utility for Blue. Increasing L would serve to reduce the force of attack, and to increase the willingness to protect.

4.2 One Attacker, N Defenders

Considering Eqs. (1 and 2), then the value of H_e is the same for all defenders in a tightly coupled network. In this case, $Red = akL(1-H_e)-(1-(1-p_c)^k)F$. With respect to k, incentives to increase the force of attack are linear and enforcement is asymptotic. The second derivative is strictly positive; maxima can only occur on the endpoints, $k \in \{0, N\}$. Intuitively, an attacker who does not want to leave "cash on the table" will either attack all defenders (rather than a subgroup) or will remain passive.

Internal equilibria, if they exist, are of the general form (with H_0 being the contribution function if the defender defects to passivity or self-insurance unilaterally):

$$H_e = 1 - \frac{(1-(1-p_c)^N)F}{aNL} \tag{15}$$

$$a = \frac{b}{L}(H_e - H_0)^{-1} \tag{16}$$

In the following we investigate the five different canonical contribution functions to identify Nash equilibria. Note that buying both protection and self-insurance at the same time is strictly dominated for nonzero b and c in all scenarios. If not indicated otherwise all defender strategies are symmetric (i.e., all Blue will select the same strategy).

Total Effort: In a total effort game, $H_e = \frac{1}{N}\sum_{i=1}^{N} e_i$. The second derivative test indicates that the optimal strategies must be corner cases. The conditions to select between the three strategies are as follows:

Full Protection If $Nb = \min(aL, Nb, c)$, then Blue plays (e,s)=(1,0).

Full self-insurance If $c = \min(aL, Nb, c)$, then Blue plays (0,1).

Passivity If $aL = \min(aL, Nb, c)$, then Blue plays (0,0).

Result 2: In a multiple defender total effort game, the relative importance of the cost of protection for the deterrence equilibrium decreases as N increases. Red's utility grows with N in equilibria where she is active.

Weakest-Link: In a weakest-link game, $H_e = \min(e_i)$. The second derivative test indicates that self-insurance is monotone, but protection may have an internal maximum. Therefore, the pure strategies are of the form $(e_i, s_i) \in \{(0,0),(0,1),(\hat{e}_0,0),(\hat{e}_0,1)\}$ with $\hat{e}_0$ being a uniformly chosen protection effort of all players. Since $(\hat{e}_0, 1)$ is dominated the conditions for Nash equilibria are as follows:

Protection. If $aL > b$ and $\hat{e}_0 > \frac{aL-c}{aL-b}$, then Blue may coordinate on $(\hat{e}_0, 0)$ for any $\hat{e}_0$ between $\frac{aL-c}{aL-b}$ and an upper boundary value. For an exogenous non-strategic attacker the upper boundary is 1 [22]. Considering a strategic attacker we find that interior solutions with $(0 \leq \hat{e}_0 \leq 1)$ and $(0 \leq a \leq 1; k = N)$ may exist. Further, when the upper boundary is less than 1 (conditions can be determined from Eqs. 15 and 16) the threat of high protection may discourage the attacker but also lower the incentives for defenders to invest in protection.

Full Self-insurance. If $c = \min(aL, aL(1-\hat{e}_0) - b\hat{e}_0, c)$, then Blue plays (0,1).

Passivity. If $aL = \min(aL, b, c)$, then Blue plays (0,0).

Result 3: In the case that full self-insurance costs more than the expected losses with protection, Red's decisions are identical to her choices in the one-on-one game and she attacks all possible targets. On the other hand, if there is a chance that the defenders would have to settle for a low $\hat{e}_0$ and full self-insurance costs less than the expected losses with this protection level then Blue can profitably defect to a self-insurance strategy. Therefore, the ability to coordinate on a high $\hat{e}_0$ is extremely important to defenders.

Protection equilibria become increasingly unlikely with increasing N if we assume that there is at least a small chance that each individual fails to coordinate successfully on a common protection level [22,40]. As Varian suggests "weakest link technology confers an advantage to small [defender] teams" [41]. Red benefits from such coordination failures.

Best Shot: In a best shot game, $H_e = \max(e_i)$. As shown in [22], there is no case in a best shot game with homogeneous defenders in which all defenders choose protection. This is easy to show with an indirect proof: If we assume there is a protection equilibrium for non-trivial parameters, then any single Blue player could profitably deviate by free-riding on his teammates [41]. Because of this, there is no symmetric pure protection equilibrium. Increasing the number of players has no effect on this finding.

Result 4: Due to an inability to coordinate on protection, defenders will prefer to shirk on protection and are vulnerable to a motivated attacker. With $b > c$ defenders will select full self-insurance. If both costs are larger than the expected loss defenders will remain passive.

k-Weakest-Target Game without mitigation: In the following we consider games for loosely coupled contribution functions. Let $\hat{e}$ = the k-th smallest e chosen by any defender i. Any Blue player choosing $e > \hat{e}$ would switch to $\hat{e} + \eta$, where $\eta \to 0$. In that case every player choosing $e < \hat{e}$ would choose $\hat{e} + 2\eta$, thus destabilizing any pure protection strategy attempts with a non-strategic attacker [22]. In Appendix A we include the detailed derivations for a mixed strategy equilibrium. Below we summarize the results.

We can derive the probability distribution function of self-protection in a mixed Nash equilibrium:

$$f = \frac{f_{e*}}{(1 + 2(2k - N)f_{e*}(e - e*))} \tag{17}$$

$$\text{where } f_{e*} = \frac{b}{aL(N-k)\binom{N-1}{j}} \tag{18}$$

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

$$q = .5 + (\sum_{j=0}^{k-1} \binom{N-1}{j} - \frac{c}{aL} 2^{N-1}) / \binom{N-1}{k-1} (N - k) \tag{19}$$

Result 5: If k is not limited, Red will always play (1,N). A mixed strategy for defenders exists. The defensive strategy is given by Eqs. (17 - 19).

k-Weakest-Target Game with Mitigation: A more nuanced version of the above game allows players a degree of individual protection in a loosely coupled scenario. In this case, a pure full protection equilibrium exists as long as protection is less expensive than self-insurance. Furthermore, to find additional mixed strategies an analysis quite similar to the above can find a probability distribution of strategies for Blue. Please refer to Appendix A for the general approach to derive the results. The probability distribution function f of self-protection in a mixed Nash equilibrium is:

$$f = \frac{\frac{b}{aL} - .5^{N-1} \sum_{j=0}^{k-1} \binom{N-1}{j} + \binom{N-1}{k-1}(N-k).5^{N-2} f_{e^*}(e-e^*)}{(1-e)\binom{N-1}{k-1}(N-k).5^{N-2}[1+2(2k-N)f_{e^*}(e-e^*)]}$$

$$\text{where } f_{e^*} \approx [\frac{b}{aL} - .5^{N-1} \sum_{j=0}^{k-1} \binom{N-1}{j}]/(1-e^*)\binom{N-1}{k-1}(N-k).5^{N-2}$$

This distribution is asymptotic at $e = 1$, indicating the benefit of mitigation. Interestingly, the probability of self-insurance is identical to the unmitigated case (see Eq. 19).

From Red's point of view, k is no longer necessarily increasing after its second root. Increasing k too high will force Blue to protect. In this case, because Red is monotone in a, she can first maximize this parameter. She will then choose k such that the cumulative binomial distribution is smaller than the cost benefit ratio, $(k, N, e^*) < b/L$. Blue then backs down into the mixed strategy, leading to a Nash equilibrium.

Result 6: In the weakest target game with mitigation we find that Red actually attacks fewer targets (but with more force) compared to the other games, and Blue players protect and self-insure according to their mixed strategy. Furthermore, as N increases, so does the number of targets that Red attacks.

4.3 M Attackers, N Defenders

Now that the various forms of contribution functions have been analyzed we can generalize from one attacker to $M \in \mathbb{N}$ attackers. We denote with m $(0 \leq m \leq M)$ the number of players who decide to engage in offensive actions. Assuming that Blue does not suffer multiple losses from being compromised by one attacker or many, we find Red's new attack force, a_j, by substitution.

Let a be the total strength of all attackers, and a_j the strength of an individual, we can substitute $(1 - (1 - a_j)^m) = a$ into Eq. 11. That is, we assume that defenders suffer from an increased attack force when multiple malefactors engage in offensive actions. Rearranging we find the new strategy, $a_j = 1 - (1 - \frac{b}{L})^{1/m}$. As the number of attackers, m, increases (given a fixed number of defenders, N), each Red will attack with proportionally less force in every game where Red plays an interior strategy. Given a sufficiently large increase in the number of attackers, the resulting decrease in attack force necessary for an interior outcome creates disincentives for attackers to be active considering the expected fine. At this tipping point the group of attackers is deterred from attacking simultaneously. However, if all the Red quit attacking at once, then it becomes profitable for an individual malefactor to restart her offensive efforts, resulting in an unstable outcome. As the number of attackers grows large, they begin to suffer coordination problems (similar to defenders in the best shot game).

Result 7: For tightly coupled games, we can derive the tipping point as m increases (with a being the total aggregate strength of all attackers):

$$(1-(1-a)^{1/m})NL > p_cF \tag{20}$$

$$m > \frac{ln(1-a)}{ln(1-\frac{p_cF}{NL})} \tag{21}$$

This finding could explain several practices observable with modern malware. For example, security researchers have recorded special cases where worms are coded to attack and replace other worms (e.g., the Netsky email virus removed Mydoom infections), or to strengthen the defenses of a compromised machine to prevent the infiltration by other malicious code (e.g., by downloading patches). Some malware authors utilize command-and-control infrastructures that allow them to throttle attacks, limit damages to compromised machine that might get users' attention (e.g., popups) and, more generally, avoid saturation effects.

5 Conclusions

There are several key findings from this research:

Nash Equilibria: Although the boundaries vary, these games all share common classes of Nash Equilibria:

- Full Attack: In the case that either the cost of self-insurance or the maximum loss is strictly less than the cost of protection, Red attacks with full force, and Blue suffers that cost or self-insures as appropriate.
- Deterrence: If the fine is so high that attacking with any force is not profitable, Red will not attack at all, and Blue need not protect or self-insure.
- Interior Equilibria: There are certain games (as in the weakest-link) where the attacker is active, and the defender protects, but not fully.

Non-equilibrium states: There are several states where pure equilibria do not exist. First, the weakest target game without mitigation and the best shot game do not offer pure symmetric protection strategies. Second, if the number of attackers increases, the network might reach a state of saturation creating coordination problems for the attackers.

Attackers: Including attackers in the game-theoretic model has several important implications. For example, expanding to the multiplayer case, there is an asymmetry between attackers and defenders. Because attackers can attack multiple targets, they can attack fewer defenders and still be profitable. This pushes defenders into undesirable states of protecting when attackers do not attack or not protecting when they do. Taking into account strategic attackers, full protection equilibria become increasingly unlikely.

Loosely and Tightly Coupled Contribution Functions: The attacker's strategy depends on the nature of the contribution function just as much as this is the case for defenders. On the one hand, in the case of a tightly coupled contribution function, attacking all defenders strictly dominates attacking a subset. On the other hand, this is not necessarily true in a loosely coupled game. Instead, it may be more profitable to target fewer defenders, but with more force.

Deterrence: Attackers may be deterred from attacking if the expected fine outweighs the expected earnings from an attack. This occurs when the attacker's break even point is greater than N. In other words, there are not enough targets to be profitable. This does imply that a government could attempt to set enforcement levels and fines such that attackers will be deterred.

Asymmetry: The fact that Red can attack many targets leads to an asymmetrical game where Red has more ability to control the state of the game than Blue.

Attacker Coordination: Bounded attacks become less likely as the number of attackers increases. If the attackers are not coordinated, they will eventually attack with too much force causing the defenders to protect. Compared to a deterrence equilibrium, this is costly for both the defenders and the attackers. This implies that sophisticated attackers will rely on command-and-control infrastructures rather than autonomous agents to manage the spread of their code. These findings also suggest that malware authors will attempt to make their code appear sufficiently benign, so that defenders are not incentivized to protect against it.

Another way that attackers may solve the coordination problem is through the open market. Phishers started to develop a market economy in which also botnet herders participate [1,16]. Botnets can now be rented for spam campaigns and distributed denial of service (DDoS) attacks [42]. This kind of marketplace could have several effects: by leasing time on their bots attackers get additional utility; by utilizing a market it may become harder to track who really launched an attack, decreasing the likelihood of being caught; and this process also significantly reduces the barrier to entry for launching distributed attacks.

Limitations and future work: We have made several assumptions, for example, the homogeneity of the players. In prior work, we have shown that heterogeneity can have a significant impact on defenders' strategies [23]. Other assumptions include the perfect attack and defense assumptions. In reality, there is often no such thing as either. As Anderson points out [3], there is often an asymmetry in finding exploits that favors the attacker.

We have not explicitly accounted for research and reconnaissance costs. These would serve as a barrier to entry for potential attackers. Furthermore, we have assumed that attackers are not directly turning against each other. In reality, rival botnets may be more tempting targets than 'civilians,' and botnet hijacking has been observed 'in the wild' [24].

Another limitation is the assumption of symmetry between the loss for defender and the gain for attackers. We can consider divergent subjective utilities: a) the defense loss is higher (then we would expect deterrence equilibria to be more common), or b) the offense gain is higher (then we would expect internal equilibria to be most common). Similarly, it may not be always the case that an attacker will benefit from a security compromise if the defender is self-insured. For example, installing spyware to gather personal information is of reduced utility if the defender has implemented a credit alert or freeze.

Possible extensions include a model of defensive hacking and activities of vigilante defenders [32]. There are significant economic and ethical questions when defenders can counterattack. If a vigilante defender compromises a botnet, and damages an infected machine, it may be for the greater good, but there is a personal risk of legal liability. This is further complicated by the fact that computer security has become highly industrialized [34]. Firms providing security services and research may be in the best position to actually implement vigilante hacking. But simply eliminating attackers would reduce the need for their products.

The present analysis relies on game theory and, in particular, Nash equilibrium analysis. We plan to expand the analysis to different behavioral assumptions to narrow the gap between formal analysis and empirical observations in the field and the laboratory [21].[6] Notwithstanding, we expect that the results provided in this paper will be of interest to security practitioners and researchers alike.

References

1. Abad, C.: The economics of phishing: A survey of the operations of the phishing market. First Monday 10(9) (2005)
2. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. IEEE Security & Privacy 3(1), 26–33 (2005)
3. Anderson, R.: Why information security is hard - an economic perspective. In: Proceedings of the 17th Annual Computer Security Applications Conference (ACSAC 2001), New Orleans, LA (December 2001)
4. Aspnes, J., Chang, K., Yampolskiy, A.: Inoculation strategies for victims of viruses and the sum-of-squares partition problem. Journal of Computer and System Sciences 72(6), 1077–1093 (2006)
5. Becker, G.: Crime and punishment: An economic approach. Journal of Political Economy 76(2), 169–217 (1968)
6. Böhme, R., Kataria, G.: Models and measures for correlation in cyber-insurance. In: Proceedings of the Fifth Annual Workshop on Economics and Information Security (WEIS 2006), Cambridge, UK (June 2006)

[6] See, for example, the application of near rationality to different network games [11].

7. Bolot, J., Lelarge, M.: A new perspective on internet security using insurance. In: Proceedings of the 27th Conference on Computer Communications (INFOCOM 2008), Phoenix, AZ, April 2008, pp. 1948–1956 (2008)
8. Cameron, S.: The economics of crime deterrence: A survey of theory and evidence. Kyklos 41(2), 301–323 (1988)
9. Cavusoglu, H., Raghunathan, S., Yue, W.: Decision-theoretic and game-theoretic approaches to IT security investment. Journal of Management Information Systems 25(2), 281–304 (Fall 2008)
10. Chantler, N.: Profile of a Computer Hacker. Interpact Press, Seminole (1997)
11. Christin, N., Grossklags, J., Chuang, J.: Near rationality and competitive equilibria in networked systems. In: Proceedings of ACM SIGCOMM 2004 Workshop on Practice and Theory of Incentives in Networked Systems (PINS), Portland, OR, August 2004, pp. 213–219 (2004)
12. Chung, W., Chen, H., Chang, W., Chou, S.: Fighting cybercrime: a review and the taiwan experience. Decision Support Systems 41(3), 669–682 (2006)
13. Clark, D., Konrad, K.: Asymmetric conflict: Weakest link against best shot. Journal of Conflict Resolution 51(3), 457–469 (2007)
14. Cornes, R., Sandler, T.: The theory of externalities, public goods, and club goods, 2nd edn. Cambridge University Press, Cambridge (1996)
15. Cremonini, M., Nizovtsev, D.: Understanding and influencing attackers decisions: Implications for security investment strategies. In: Proceedings of the Fifth Annual Workshop on Economics and Information Security (WEIS 2006), Cambridge, UK (June 2006)
16. Cymru, T.: The underground economy: Priceless. ;login: The USENIX Magazine 31(6) (2006)
17. Danezis, G., Anderson, R.: The economics of resisting censorship. IEEE Security & Privacy 3(1), 45–50 (2005)
18. Franklin, J., Paxson, V., Perrig, A., Savage, S.: An inquiry into the nature and causes of the wealth of internet miscreants. In: Proceedings of the 14th ACM Conference on Computer and Communications Security (CCS 2007), Alexandria, VA, October/November 2007, pp. 375–388 (2007)
19. Gordon, S.: The generic virus writer. In: Proceedings of the International Virus Bulletin Conference, Jersey, Channel Islands, September 1994, pp. 121–138 (1994)
20. Granick, J.: Faking it: Calculating loss in computer crime sentencing. I/S: A Journal of Law and Policy for the Information Society 2(2), 207–228 (Spring/Summer 2006)
21. Grossklags, J., Christin, N., Chuang, J.: Predicted and observed behavior in the weakest-link security game. In: Proceedings of the USENIX Workshop on Usability, Privacy and Security (UPSEC 2008), San Francisco, CA (April 2008)
22. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: Proceedings of the 2008 World Wide Web Conference (WWW 2008), Beijing, China, April 2008, pp. 209–218 (2008)
23. Grossklags, J., Christin, N., Chuang, J.: Security and insurance management in networks with heterogeneous agents. In: Proceedings of the Ninth ACM Conference on Electronic Commerce (EC 2008), Chicago, IL, July 2008, pp. 160–169 (2008)
24. Higgens, K.J.: Dark Reading (April 2007)
25. Hirshleifer, J.: From weakest-link to best-shot: the voluntary provision of public goods. Public Choice 41(3), 371–386 (1983)

26. Jiang, L., Anantharam, V., Walrand, J.: Efficiency of selfish investments in network security. In: Proceedings of the 2008 Workshop on the Economics of Networks, Systems, and Computation (NetEcon 2008), Seattle, WA, August 2008, pp. 31–36 (2008)
27. Kesan, J., Majuca, R., Yurcik, W.: Three economic arguments for cyberinsurance. In: Chander, A., Gelman, L., Radin, M. (eds.) Securing Privacy in the Internet Age, pp. 345–366. Stanford University Press, Stanford (2008)
28. Kessler, G.: Defenses against distributed denial of service attacks (2000)
29. Kshetri, N.: The simple economics of cybercrimes. IEEE Security & Privacy 4(1), 33–39 (2006)
30. Kunreuther, H., Heal, G.: Interdependent security. Journal of Risk and Uncertainty 26(2–3), 231–249 (2003)
31. Moscibroda, T., Schmid, S., Wattenhofer, R.: 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 2006), Denver, CO, July 2006, pp. 35–44 (2006)
32. Naraine, R.: Kraken botnet infiltration triggers ethics debate. eWeek.com (May 2008)
33. Pautasso, C., Wilde, E.: Why is the web loosely coupled? A multi-faceted metric for service design. In: Proceedings of the 2009 World Wide Web Conference (WWW 2009), Madrid, Spain, April 2009, pp. 911–920 (2009)
34. Potter, B.: Dirty secrets of the security industry. Defcon XV, Las Vegas (2007)
35. Radosavac, S., Kempf, J., Kozat, U.: Using insurance to increase internet security. In: Proceedings of the 2008 Workshop on the Economics of Networks, Systems, and Computation (NetEcon 2008), Seattle, WA, August 2008, pp. 43–48 (2008)
36. Schechter, S., Smith, M.: How much security is enough to stop a thief? In: Wright, R.N. (ed.) FC 2003. LNCS, vol. 2742, pp. 122–137. Springer, Heidelberg (2003)
37. Schneier, B.: Tactics, targets, and objectives. Wired.com (May 2007)
38. Swire, P.: No Cop on the Beat: Underenforcement in E-Commerce and Cybercrime. Journal on Telecommunications and High Technology Law, forthcoming (2008)
39. The Honeynet Project. Know your enemy: the tools and methodologies of the script-kiddie (July 2000), http://project.honeynet.org/papers/enemy/
40. Van Huyck, J., Battallio, R., Beil, R.: Tacit coordination games, strategic uncertainty, and coordination failure. American Economic Review 80(1), 234–248 (1990)
41. Varian, H.: System reliability and free riding. In: Camp, L., Lewis, S. (eds.) Economics of Information Security. Advances in Information Security, vol. 12, pp. 1–15. Kluwer Academic Publishers, Dordrecht (2004)
42. Weinberg, N.: Botnet economy runs wild. Network World (April 2008)

A Mixed Strategy for Weakest Target Game Without Mitigation

We investigate whether a mixed strategy can be derived. Assume there is a cumulative distribution of protection strategies F. We can use the cumulative distribution of the binomial distribution to represent the chance that a player will be compromised given a fixed e. The expected utility of Blue is:

$$Blue = aL\sum_{j=0}^{k-1} \tbinom{N-1}{j} F_e^j (1-F_e)^{N-1-j} - be_i - cs_i \tag{22}$$

In Nash equilibria, the first order condition must hold:

$$\begin{aligned} 0 &= aL(N-k)\binom{N-1}{j}F_e^{k-1}(1-F_e)^{N-1-k}(f) - b \\ \frac{b}{aL(N-k)\binom{N-1}{j}} &= F_e^{k-1}(1-F_e)^{N-1-k}(f) \\ \frac{b}{aL(N-k)\binom{N-1}{j}} &= exp\{(k-1)lnF_e + (N-1-k)ln(1-F_e)\}(f) \\ f &= \frac{b}{aL(N-k)\binom{N-1}{j}exp\{(k-1)lnF_e + (N-1-k)ln(1-F_e)\}} \end{aligned}$$

Then we can expand the exponentiated part about e^* = the median of f using a Taylor expansion. Thus,

$$f = \frac{b}{aL(N-k)\binom{N-1}{j}(\frac{1}{2})^{N-2}(1+2(2k-N)f_{e^*}(e-e*))} \tag{23}$$

$$\text{where } f_{e^*} = \frac{b}{aL(N-k)\binom{N-1}{j}} \tag{24}$$

$$\text{thus } f = \frac{f_{e^*}}{(1+2(2k-N)f_{e^*}(e-e*))} \tag{25}$$

The approximation of f about e^* is asymptotic as $e \rightarrow e*$. Knowing that Blue will never play $e > aL/b$ because of dominance, we estimate $e^* = aL/b$.

If insurance is not overpriced, then we know $F(0) = q; Blue(0,0) = c$:

$$pl\sum_{j=0}^{k-1}\binom{N-1}{j}q^j(1-q)^{N-1-j} = c \tag{26}$$

Using a Taylor expansion again, we find:

$$\frac{1}{2})^{N-1}\sum_{j=0}^{k-1}\binom{N-1}{j} - (\frac{1}{2})^{N-1}\binom{N-1}{k-1}(N-k)(q-.5) = c/aL \tag{27}$$

$$-\binom{N-1}{k-1}(N-k)(q-.5) = \frac{c}{aL}2^{N-1} - \sum_{j=0}^{k-1}\binom{N-1}{j} \tag{28}$$

$$q = .5 + (\sum_{j=0}^{k-1}\binom{N-1}{j} - \frac{c}{aL}2^{N-1})/\binom{N-1}{k-1}(N-k) \tag{29}$$