Distributed attacks as security games

Neal Fultz
School of Information
UC Berkeley
nfultz@ischool.berkeley.edu

ABSTRACT

Due to the development of easy-to-use software, distributed denial of service attacks have gone from theoretical weapons to being in the hands of criminals in less than twenty years. This paper analyzes the threat of DDoS attacks by developing a two-sided multiplayer model of security in which attackers attempt to deny service and defenders attempt to protect against attacks. Key findings include the existence of protection and non-attack Nash equilibria among different configurations of contribution functions and numbers of players, and validation of this model against infection data. I find that strategic attackers launch attacks only if protection will not occur. Therefore, the threat of protection can be enough to deter an attacker, but as the number of attackers grows, this equilibrium becomes increasingly unstable.

Categories and Subject Descriptors

C.2 [**Computer Systems Organization**]: Computer Communication Networks; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; K.4.4 [**Computers and Society**]: Electronic Commerce—*Security*

General Terms

Economics,Security

Keywords

Game Theory, Economics of Security, Distributed Denial of Service Attacks, Botnets

1. INTRODUCTION

In April 2007, Estonian government and financial institutions were attacked in what has been called the first cyberwar. Critical systems, such as emergency telephone numbers, were compromised. The attack came over the internet, and was initiated by one person [19] [1].

Distributed denial of service attacks themselves are not technically sophisticated. All that is necessary is enough clients using enough resources to overwhelm a server. Perhaps the least sophisticated, online 'sit-ins' date back to at least 1995 [6]. Since then, computing power and bandwidth have followed Moore's law and grown exponentially; it now takes more than a few hundred volunteers to knock out a website.

On the other hand, attackers have become increasingly sophisticated. Computer exploits are cataloged and traded, backdoors allow attackers to remote control a computer, and rootkits let them mask their presence. Viruses and worms can compromise millions of computers quickly. Some computers don't even need to be compromised to be turned into a weapon [15]. In the case of Estonia, a teenager used a network of compromised machines (a "botnet") to shut down government services for nearly two weeks.

Botnets are how attackers are keeping up in the malware arms race today. A machine can be infected from tainted email, or from receiving a worm. At that point, the compromised computer ("zombie") connects to a command-and-control channel, such as IRC or a website, and awaits commands from its new master. Typically, these zombies are used to send spam, host pirated software, and launch DDoS attacks.

Botnets have made distributed denial of service attacks easy, and these are becoming increasingly common for both political and profit motives [18]. Although they were first used by hackers to show off technical prowess by shutting down big name websites [11], recently they have been used as political weapons in Tibet by protesters [7], and also as extortion tools by criminals [12].

There are important economic questions surrounding these attacks. How do attackers benefit from these attacks, and how are the victims hurt? Why do the attackers attack? Why do the victims choose not to protect themselves? If we could have changed a variable, how would that have affected the attack?

2. BACKGROUND AND MOTIVATION

The economics of security is still an emerging field [3]. It has been said that people are the eighth network layer; that is a strong reason to apply the lens of economics to this domain. Examples of misaligned incentive structures in the security and privacy field are numerous. The tragedy of the commons, the market for lemons, and other classic economic ideas all have applications to the field of security.

Research on formalized models followed as the security and economics communities began to interact. Varian treats security as a public good under a static threat, and applies classic public goods models to much effect. His models are very general; this could be construed as a strength or a weakness [20]. Perhaps most interestingly, he briefly considers adversarial games. Grossklags expands on these models by introducing a private security variable; this results in total security becoming a hybrid good. He also bounds single player exertions, which leads to significantly different results [8].

On the other side of the DDoS, Schechter draws upon economics of crime literature to construct a model of attackers, and derives what kinds of penalties and probabilities of enforcement will deter an attacker [17]. The basic model of crime economics is that criminals are utility optimizers who evaluate the risk and rewards of committing a crime[4]. Schechter puts these kinds of models specifically in a computer security context.

Also relevant is the study of computer viruses. Prior work has successfully shown that epidemiological models can be applied to computer viruses and produce good descriptive models [23]. Other work has been done to cut through inflated reporting to try to discern the true costs and the worst case costs of viruses [21]. Botnets have been more difficult to model, but

This paper attempts to combine the adversarial aspect of of Varian, the hybrid security good aspect of Grossklags, and Schechter's deterrence model and adapt this to the context of wide-scale distributed attacks instead highly targeted small scale attacks. Further, we can then fit this model to data from DDoS attacks, and estimate the true costs of protection and the attacker's expected probability of being caught.

3. ONE SHOT GAMES

Security games are defined as games that model the decision making process of network security in which members may purchase security against a risk. The net effect of that purchase is determined by a contribution function, which may take into account other players' purchases. Depending on the functional form of the contribution function, security may be a public good , a private good, or a hybrid of the two [20] [8].

However, the most important factor to modeling a DDoS is the fact that there are multiple defenders and multiple attackers.

Building on Grossklags et all's analysis, I consider the five canonical contribution functions of security games, but have generalized from one-shot games to repeated interactions, and also from weakest-target to k-weakest-targets. Perhaps more significantly, I also account for the possibility of attackers capable of acting strategically.

Assumptions. These models are all built on several underlying assumptions: (i) the contribution function is shared by all defenders, (ii) each player acts independently and simultaneously, (iii) costs are identical for all defenders. (iv) the attacker has available a perfect attack (v) the defender has available a perfect defense (vi) the perfect defense trumps the perfect attack.

From these assumptions, a formalized model is constructed; each of $N \in \mathbb{N}$ defenders faces a maximum loss $L > 0$ if they are attacked and compromised by at least one attacker. Defenders may mitigate this threat by purchasing e_i percent protection at a cost of b per unit, or s_i percent self-insurance at a cost of c per unit. Thus, $0 \leq e_i \leq 1$ and $0 \leq s_i \leq 1 \forall i \in N$. Self insurance lowers losses deterministically, whereas protection lowers losses probabilistically through the contribution function.

Attacks arrive with a probability p_a, chosen by $M \in \mathbb{N}$ attackers. Attackers face a fine $F > 0$ which they suffer only if they are caught with a probability p_c. p_c is assumed to be independent of $p_a \forall p_a > 0$. If an attack is successful, the attacker receives L from each defender they compromise. In the weakest targets case, the attacker may also choose a value $k \in \mathbb{N}, 0 \leq k \leq N$ of how many defenders they attack.

Given the contribution function H_e, the expected utility of defender i is:

$$E(U_i) = -(1-(1-p_a)^m)L(1-H_e)(1-s_i) - be_i - cs_i \quad (1)$$

Similarly, the expected utility of attacker i is:

$$E(U_i) = \sum_{1}^{k} p_a L(1-H_e) - (1-(1-p_c)^k)F \quad (2)$$

For the purposes of simplifying notation, let $Blue_i = E(U(i)$ for defenders, $Red_i = E(U_i)$ for attackers,

3.1 One Attacker, One Defender

In the case of one attacker and one defender, assuming $H_e = e_1, k = 1$, the utility functions simplify to:

$$Blue_i = -p_a L(1-e) - be - cs \quad (3)$$

$$Red_i = p_a L(1-e) - p_c F \quad (4)$$

First, consider Blue. The second derivative test indicates that Blue is monotone in both e and s. Therefore, the maximum of Blue can only exist in the corner cases of $(e,s) \in \{(0,0),(0,1),(1,0),(1,1)\}$. Substituting, we find Blue equals $-p_a L, -b, -c, and -(b+c)$, respectively. Dominance implies that $(1,1)$ will not be played if both costs are positive. We find the following three strategies:

Passivity. If $p_a L = min(p_a L, c, b)$, Blue plays (0,0).

Full Insurance. If $c = min(p_a L, c, b)$, Blue plays (0,1).

Full Protection. If $b = min(p_a L, c, b)$, Blue plays (1,0).

Second, consider Red's best reply to each strategy. First, Red will not attack if $p_c F > L$.

If $L = min(L, c, b)$, the Blue will play passivity $\forall p_a$. Red's best reply is $p_a = 1$. As neither player can profitably deviate, this is Nash.

Similarly, if $c = min(L, c, b)$, the Blue will play either full insurance $\forall p_a L > c$ and passivity $\forall p_a L \leq c$. Insurance has no effect on Red, so Red's best reply is still $p_a = 1$, and these strategies are also Nash.

Finally, consider $b = min(L, c, b)$. Blue plays passivity if $p_a L \leq b$, and full protection otherwise. Red risks the fine with no possible reward if he plays p_a such that Blue protects, and risks the fine with a diminished reward if he plays $p_a = b/L$. If this diminished payout is non-negative, Red will attack. Therefore, Red plays $(p_a, k) = (b/L, 1)$ if $b > p_c F$ and (0,0) otherwise. This is also Nash.

Results. In all three cases, Blue does not protect. Interestingly, in the case where protection is not overpriced relative to losses and self-insurance, the cost-benefit ratio of full protection serves as an upper bound on Red's p_a. Therefore, reducing b would lead to less frequent attacks and a higher expected value for Blue. Increasing L would also serve to reduce the frequency of attack, but would not increases Blue's expectation. In the event that $b \leq p_c F$, the availability of cheap protection serves as a deterrent to Red.

3.2 One Attacker, N Defenders

$$Blue = -p_a L(1 - H_e)(1 - s_i) - be_i - cs_i \tag{5}$$

$$Red = \sum_{1}^{k} p_a L(1 - H_e) - (1 - (1 - p_c)^k)F \tag{6}$$

In a tightly coupled network, $H_{(e)}$ evaluates to the same value for all players; the vector of inputs is being reduced to a single value. In this case, $Red = p_a k L(1 - H_{(e)}) - (1 - (1 - p_c)^k)F$. Over k, Red is convex and negative until some root r given p_a. Therefore, if r > N, then it is never profitable for Red to attack, so k=0 (no attack) strictly dominates all other k. If on the other hand $r \leq N$, playing $k = N$ strictly dominates all other k. Let us assume it is possibly profitable for Red to attack and investigate the canonical contribution functions.

Total Effort:. In a total effort game, $H_e = \frac{1}{N}\sum_{i=1}^{N} e_i$. Once again, the second derivative test indicates that Blue is monotone over s and i, so the maximal strategies must be the corner cases:

Passivity. If $p_a L = min(p_a L, Nb, c)$, Blue plays (0,0).

Full Insurance. If $c = min(p_a L, Nb, c)$, Blue plays (0,1).

Full Protection. If $Nb = min(p_a L, Nb, c)$, Blue plays (1,0).

As in the 1-on-1 game, if Blue plays passivity or insurance, Red's best reply is full attack (1,N). If Blue plays protection, Red's best reply is either to not attack at all if $Nb/L < P_c F$, or to play $(p_a, k) = (Nb/L, N)$ otherwise. These strategies can be shown to be Nash.

Results. In a multiplayer Total Effort game, the deterrent effect of protection decreases as N increases. In a scenario with where protection is possible, Red's optimal probability of attack grows with N, and his expected utility is proportional to N^2.

Weakest Link. In a weakest link game, $H_e = min(e)$. The second derivative test indicates that insurance is monotone, but protection may have an internal maxima. Assuming protection is not overpriced, let $\hat{e}_0$ be the empirical minimum of all e; if Blue plays $e_i < \hat{e}_0$, it hasn't bought as much protection as possible; if Blue plays $e_i > \hat{e}_0$, then because of the contribution function it receives no additional benefit, but pays the extra cost. Therefore, the pure strategies are $(e_i, s_i) \in \{(0,0), (0,1), (\hat{e}_0, 0), (\hat{e}_0, 1)\}$. Once again, buying both protection and insurance is strictly dominated for nonzero b and c.

Passivity. If $p_a L = min(p_a L, b, c)$, Blue plays (0,0).

Full Insurance. If $c = min(p_a L, p_a L(1 - \hat{e}_0) - b\hat{e}_0, c)$, Blue plays (0,1).

Full Protection. If $p_a L > b$ and $\hat{e}_0 > \frac{p_a L - c}{p_a L - b}$, Blue plays $(\hat{e}_0, 0)$.

In this case, Red's best reply to a nonprotection strategy is $p_a = 1$. Red's best reply to protection depends on $\hat{e}_0$ and whether L is less than c. If $L(1 - \hat{e}_0) > c$, Red's best reply is $p_a = 1$, which forces Blue to switch to insurance. Otherwise, Red plays $p_a < b/L$, and Blue switches to passivity. These strategies are Nash.

Results. In the case that insurance is overpriced relative to the expected loses with protection, Red is bounded by the cost-benefit ratio just as in the one-on-one game. On the other hand, if insurance is cheap but $\hat{e}_0$ is sufficiently small, Red can actually increase his attack and force Blue into an insurance strategy. Therefore, knowledge of $\hat{e}_0$ is extremely important to Blue, just as in [8]; as N increases, protection becomes both less likely for Blue because of coordination issues, and less of a deterrent to Red, whose payoff increases with N.

Best Shot. In a best-shot game, $H_e = max(e)$. As shown in [8], there is no case in a best shot game with homogeneous defenders in which Blue chooses protection. This is easy

to show with an indirect proof: If we assume there is a protection equilibrium for nontrivial parameters, then any single Blue player could profitably deviate by free-riding on his teammates. Because of this, Red's best reply is to always play the maximum attack, and Blue chooses the cheaper of passivity and insurance. This is Nash. Increasing the number of players has no effect on this equilibrium.

k-Weakest Targets. Now consider games for loosely coupled contribution functions, that is, H_e may be different for different Blue players. A simple one is k-Weakest targets.

$$H_e = \begin{cases} 1 & \text{if } e_i > e_{(k)} \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Let $\hat{e}$ = the k-th smallest e. 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, just like the best shot game, there are no pure Nash equilibria for this game.

Mixed Strategy Equilibrium. To complement Grossklags analysis, 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. Then the expected utility of Blue is

$$Blue = p_a L \sum_{j=0}^{k-1} \binom{N-1}{j} F_e^j (1-F_e)^{N-1-j} - be_i - cs_i \quad (8)$$

In Nash equilibria, the first order condition must hold:

$$0 = p_a L(N-k)\binom{N-1}{j} F_e^{k-1}(1-F_e)^{N-1-k}(f) - b$$

$$\frac{b}{p_a L(N-k)\binom{N-1}{j}} = F_e^{k-1}(1-F_e)^{N-1-k}(f)$$

$$\frac{b}{p_a L(N-k)\binom{N-1}{j}} = exp\{(k-1)lnF_e + (N-1-k)ln(1-F_e)\}(f)$$

$$f = \frac{b}{p_a L(N-k)\binom{N-1}{j} exp\{(k-1)lnF_e + (N-1-k)ln(1-F_e)\}}$$

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

$$f = \frac{b}{p_a L(N-k)\binom{N-1}{j}(\frac{1}{2})^{N-2}(1+2(2k-N)f_{e^*}(e-e*))} \quad (9)$$

$$\text{where } f_{e^*} = \frac{b}{p_a L(N-k)\binom{N-1}{j}} \quad (10)$$

$$\text{thus } f = \frac{f_{e^*}}{(1+2(2k-N)f_{e^*}(e-e*))} \quad (11)$$

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

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

$$pl \sum_{j=0}^{k-1} \binom{N-1}{j} q^j (1-q)^{N-1-j} = c \quad (12)$$

Using a Taylor expansion again, we find:

$$(\frac{1}{2})^{N-1} \sum_{j=0}^{k-1} \binom{N-1}{j} - (\frac{1}{2})^{N-1}\binom{N-1}{k-1}(N-k)(q-.5) = c/p_a L \quad (13)$$

$$-\binom{N-1}{k-1}(N-k)(q-.5) = \frac{c}{p_a L}2^{N-1} - \sum_{j=0}^{k-1}\binom{N-1}{j} \quad (14)$$

$$q = .5 + (\sum_{j=0}^{k-1}\binom{N-1}{j} - \frac{c}{p_a L}2^{N-1})/\binom{N-1}{k-1}(N-k) \quad (15)$$

Results. On the other side, we notice that Red will always play (1,N), which forces Blue to always play a pure nonprotection strategy. This is Nash.

k Weakest Targets with Mitigation. A more nuanced version of the above game allows players a degree of individual protection in a loosely coupled scenario.

$$H_e = \begin{cases} 1 & \text{if } e_i > e_{(k)} \\ e_i & \text{otherwise} \end{cases} \quad (16)$$

In this case, a pure protection equilibrium is possible so long as protection is less expensive than insurance. Furthermore,an analysis quite similar to the above can find a probability distribution of strategies for Blue. Expanding the expected utility about the median again, we find:

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

$$\text{all over } (1-e)\binom{N-1}{k-1}(N-k).5^{N-2}[1+2(2k-N)f_{e^*}(e-e^*)]$$

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

This distribution has an asymptote at e = 1, indicating the benefit of mitigation. Interestingly, the probability of insuring given cheap insurance remains the same as the nonmitigation case.

From Red's point of view, k is no longer necessarily increasing after it's second root. Increasing k too high will force Blue to protect. In this case, because red is monotone in p_a, it will first maximize that. It will then choose k such that the cumulative binomial distribution of $(k, N, e^*) < b/L$. Blue then backs down into the mixed strategy, leading to a Nash equilibrium. Red's payout is further reduced by the mitigating factor.

Results. In the mitigated weakest target game, Red actually attacks fewer targets more often than the other games, and Blue plays randomly according to a mixed strategy. Furthermore, as N increases, so does the number of targets that Red attacks.

3.3 M Attackers, N Defenders

Now that the various forms of contribution functions have been analyzed, generalizing from one attacker to M is fairly straightforward. Assuming that Blue does not receive an additional loss from being compromised by one Red or many, we find for Blue the new probability of being attacked by substitution, and inverting to find the new ceilings for Red based on the total ceiling of the probability of attack p_A derived earlier. In the case of loosely coupled contribution functions, let p_A be the ceiling on the cumulative distribution of the binomial distribution given f_e and N instead.

$$(1 - (1 - p_a)^m) = p_A \tag{17}$$

$$p_a = 1 - (1 - P_A)^{1/m} \tag{18}$$

This implies that as m increases, each Red will attack proportionally less in every game where Red is bounded by protection states and the contribution is tight. As m grows even larger, p_a may shrink small enough to the point that the second root of Red increases past N, deterring all Red from attacking simultaneously. However, if all the Red quit attacking at once, then it becomes profitable for an individual Red to start attacking again, and there is no Nash. In these cases, it appears that as Red gets large, they begin to suffer from coordination problems just like Blue. In tightly coupled games, this tipping point as m increases occurs at:

$$(1 - (1 - p_A)^{1/m})NL > p_c F \tag{19}$$

$$m > \frac{ln(1 - p_A)}{ln(1 - \frac{p_c F}{NL})} \tag{20}$$

This finding could explain the modern development of botnets. The population of autonomous malware eventually grows too large,and forces Blue to protect. Botnets, on the other hand, are characterized by command-and-control communication. This communication can allow their network to grow beyond this upper bound. Similarly, groups of attackers will pool their resources and coordinate their attacks rather than attacking independently.

4. REPEATED GAMES

Now I relax the assumption that these strategic interactions are played only once. To do so, I need to assert a few new ones: (i) Costs are fixed over time (ii) Players have an inherent discount factor, δ (iii) Games are played infinitely.

In this section, I do not address cheap insurance, as it has no effect on deterring Red.

Total Effort. Consider the Blue strategy "Always protect if the expected losses are higher than the cost, and noone has defected." I set up the game totals to find for which values of δ this strategy holds:

$$-p_A L(1 - \frac{N-1}{N}) - \frac{\delta}{1-\delta} < \frac{b}{1-\delta}$$

$$\delta > \frac{1 - \frac{bN}{p_A L}}{N+1}$$

Solving for p and negating to find Red's new ceiling,

$$p_A < \frac{bN}{L(\delta(N+1) - 1)}$$

Results. If δ is 1, then the probability of attack has fallen by a factor of N compared to the one-shot game. If the discount factor is 0, then the solution remains the same.

Weakest Link. Setting up a similar inequality,

$$\frac{1}{1-\delta}(-b) < \frac{1}{1-\delta}(-p_A L)$$

Delta is free, because protection is a stage game Nash.

Solving for p and negating to find Red's new ceiling,

$$p_A < \frac{b}{L}$$

Results. In the repeated case of Weakest Link, Red's ceiling is the exact same as in the one shot with an $\hat{e}_0 = 1$. However, the repeated nature of this game is enough to allow Blue to coordinate full protection rather than settling on an expected minimum.

Best Shot. Imagine a strategy where each round a different Blue player would take a turn purchasing full protection, as long as the expected losses were sufficiently high and no one had previously deviated.

$$\frac{1}{1-\delta^N}(-b) < \frac{1}{1-\delta}(-p_A L)$$

$$\frac{1-\delta}{1-\delta^N} > \frac{b}{p_A L}$$

$$p_A < \frac{b(1-\delta)}{L(1-\delta^N)}$$

Results. Unlike the one shot, the repeated version does bound Red's attack. More significantly, it actually gets

stronger as N increases rather than weaker. This is due to the diffusion of costs over several rounds rising from the turn taking nature of the strategy.

Weakest Target. Consider a similar strategy, where instead of taking turns to protect, k players take turns being a honeypot and play $\eta \to 0$ otherwise.

$$-2\eta b - \frac{\delta}{1-\delta} p_A L < \frac{1}{1-\delta^N}\left(\frac{(1-\delta^{k+1})p_A L - (1-\delta^N)\eta b}{1-\delta}\right)$$

Taking the limit,

$$\frac{1-\delta^N}{1-\delta^k} > \frac{b}{p_A L}$$

Red's ceiling is thus

$$p_A < \frac{b}{L}\frac{1-\delta^k}{1-\delta^N}$$

A second constraint is that

$$k < Nb/L$$

In the case of no mitigation and infinite Red strength, there is still nothing Blue can do other than passivity. However, in the mitigation case where k < N, we see that this strategy is proportionally strong to the ratio of k and N.

5. EXTENSIVE-FORM GAMES

These models still do not take into account any changes in the number of attackers or defenders, which is a major characteristic of automated malware; viruses are viruses because they spread. There are several rules we could use to try to capture contagious effects:

- Players that don't insure and are compromised become attackers for the rest of the game.
- Players that do insure and are compromised are dropped from the rest of the game.

Sketch of proof. Starting with M attackers and N defenders, and working in the limits of both parameters in a loosely coupled game, k players at most are compromised each turn. Of these, βk do not insure. Thus, in round x you would expect $M + \beta k x$ attackers and $N - kx$ defenders. In an infinite number of rounds, there exists some x such that the number of attackers crosses the point described in (20), and after that point attackers can no longer coordinate absent some external communication. In this case, there is no Nash and the probability of attack approaches 1, which would activate Blue's protection trigger. Once this happens, the game remains in that equilibria until it ends. Thus, we would expect the empirical cumulative distribution of attackers ("infection") over time to be monotone increasing but decelerating up until that point, at which point it remains level.

6. MODEL FITTING

Using the above, we can use the method of transformation to randomly generate numbers from the distributions described earlier, use simulations to generate infection and damage distributions, and fit that to actual data with numerical techniques. K-weakest targets with mitigation is the most nuanced of the models, and so I have studied it.

6.1 Code Red

Using the Code Red dataset available from CAIDA [13], we can fit to the infection CDF using numerical minimization of χ^2. Ignoring the lengthy tails, and assuming a fine of $10 million, and damages equal to $1.2 billion across 400,000 machines, and that self-insurance is over priced, we find that the expected cost of protection is $50.88 and the likelihood of catching an attacker is .48%, giving us a χ^2 of 47.4 with 38 degrees of freedom, which is just inside one standard deviation away from the expected. This shows that my model is consistent with the empirical CDF.

Looking at the scatter plot, the effect of the assumption of unlimited attacker strength is apparent; the theoretical distribution approaches it's limit much sooner than the empirical distribution. This could be because there are some constraining factors other than the number of targets limiting attacker spread.

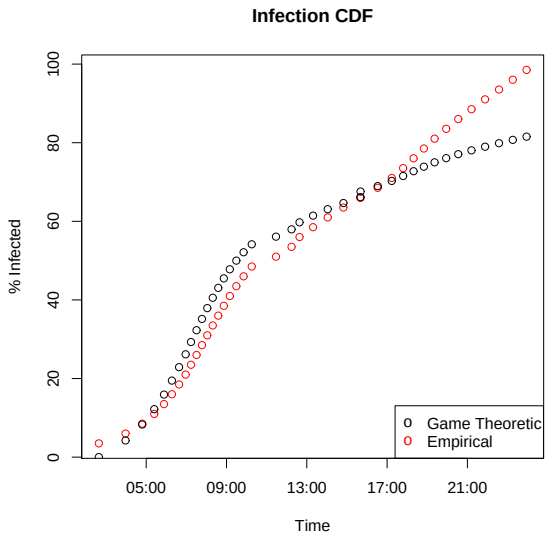

6.2 Worst Case Worm

Furthermore, we can fit to Paxon's predicted damages of a worst case worm and reusing the Code Red infection CDF [21]. Using an expected fine of $10 million, damages of $103 billion, 50 million defenders, and no cheap self insurance, we find that the expected cost of protection is approximately $20,000 and the chance of being caught is infinitesimal. This results in a minimized χ^2 of 1893197 with 38 degrees of freedom. My model clearly does not fit under these parameters.

This can be explained by the fact that either a worst-case worm will not follow the Code Red infection CDF under my model, or that a worst case attacker is not attacking strategically with respect to my model.

7. CONCLUSIONS

There are several key findings from this research:

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

- Unbounded Attack: In the case that either the cost of insurance or the maximum loss is strictly less than the cost of protection, Red attacks with full force, and Blue suffers that cost or insures as appropriate.
- Bounded Attack: If protection is cheap, Red may attack less frequently, dropping the expected loss from an attack below the cost of protection. If attacking at this level is still profitable, Red will do so, and Blue will not protect (assuming a preference of passivity over protection).
- Deterrence: If protection is so cheap that attacking at a reduced rate is not profitable, Red will not attack at all, and Blue will not protect or insure.

Non-equilibria states. There are several states where pure equilibria do not exist. These include when Blue has a preference for protection over passivity in the bounded cases, or when there are sufficient attackers that deterrence breaks down.

Attackers. Including attackers in the game theoretic model has several important implications. Attackers maximize their utility while minimizing the defense's. Expanding to the multiplayer case, there is an asymmetry between attackers and defenders. Because attackers can attack multiple targets, they can attack less frequently and still be profitable. This forces the defenders into an undesirable subgame of protecting when attackers don't attack or not protecting when they do. Taking into account attackers, protection equilibria become highly unlikely.

Loosely and Tightly Coupled Contribution Functions. Depending on the nature of the contribution function, the attacker can maximize his utility by changing different variables. In the case of a tightly coupled contribution function, the function reduces a vector of protection variables to a single value. This allows the attacker to easily determine if they should attack, by finding the second root of his utility function. If this root is less than N, they attack with k = N, and optimize over p_a. If the contribution function instead maps each element of the protection vector to a new value, instead, the expected maximum falls on $p_a = 1$, and the attacker optimizes over k instead.

Deterrence. On the other hand, attackers may be deterred from attacking at all if the expected fine outweighs the expected earnings from an attack. This occurs when the second root of the utility function is greater than N. In other words, there aren't enough targets to be profitable. However, this does imply that a government could set enforcement levels and fines such that attackers will be deterred.

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

Attacker Coordination. Deterrence becomes less likely as the number of attackers increase; if the attackers are not coordinated, eventually the attackers will over-attack, causing the defenders to protect. Compared to a deterrence equilibrium, this is costly for both the defenders and the attackers.

This implies that future attackers will rely on botnets with command and control communications rather than autonomous agents. It also implies that malware will become increasingly benign, so that defenders are not incentivized to protect against it.

A second way that attackers may solve the coordination problem is through the open market. Phishers have already developed a market economy, of which botnets are a slice [2] [5]. Although these botnets have traditionally been used for spam, they are now also rented for DDoS attacks [22]. This kind of marketplace could have several effects: by leasing time on their bots, attackers get additional utility; by going through a market, it becomes harder to track who really launched an attack, decreasing the chance of being caught; it also significantly reduces the barrier to entry for launching a DDoS attack.

Defender Coordination. There is a similar coordination problem with defenders; this paper has shown that these issues can be overcome in certain repeated cases as long as the defense has a sufficient discount factor. Defenders have a slight advantage because there are several cases where the defenders can coordinate inside the game while the attackers can not.

Limitations. In developing this model, I have made several assumptions. One major one is the homogeneity of the players. Grossklags has shown that relaxing the homogeneity assumptions can have significant impact on very similar models [9]. Other assumptions include the perfect attack and perfect defense assumptions. In reality, there is no such thing as either; as Anderson points out in [3], there is an asymmetry in finding exploits that favors the attacker, which this model does not address. Furthermore, the assumption that players act rationally is definitely questionable if the attackers are actually teenage hackers.

I have also assumed that attackers are not attacking each other. In reality, rival botnets may be more tempting targets than 'civilians,' and botnet hijacking is not unknown [10]. Malware in the wild has been observed removing other malware, which would certainly reduce damages to a defender, and over time could lead to some strange symbiotic relationship. This also brings to mind defensive hacking and vigilante defenders [14]. There are significant economic and ethical questions when defenders can counter-attack. If a vigilante defender compromises a botnet, and damages an infected machine, it may be for the greater good / social optima, but there is a personal risk of legal liability. This is

further complicated by the fact that computer security has become highly industrialized [16]. Firms providing security services and research are in the best position to actually implement vigilante hacking, but simply eliminating attackers would reduce the need for their products.

Another key limitation is the assumption symmetric utility. In the case of highly divergent subjective utilities, there are two cases: the defense is higher, and the offense is higher. If the defense's is higher, I would expect deterrence equilibria to be most common; if the offense's is higher, I would expect bounded attack equilibria to be most common.

Future Research. In the future, I would like to more formally explore the different effects arising from tightly coupled and loosely coupled contribution functions. Heterogeneous defenders and attackers both need to be analyzed, as do endgame strategies for finitely repeated games, and cases of subjective utility. Additionally, hybrid attacker models should be developed to cover vigilante defenders. Thus, the model can be generalized from DDoS attacks to a full model of cyberwar.

8. ACKNOWLEDGMENTS

Special thanks to Jens Grossklags.

Support for the CAIDA Dataset on the Code-Red Worms was provided by Cisco Systems, the US Department of Homeland Security, the National Science Foundation, DARPA, and CAIDA Members.

Drs. John Chuang and Doug Tygar.

Johnson Nguyen, Kevin Lim, Zach Gillen, Josh Gomez.

9. REFERENCES

[1] Newly nasty; cyberwarfare is becoming scarier. *The Economist*, May 2007.

[2] C. Abad. The economics of phishing: A survey of the operations of the phishing market. *First Monday*, 10(9), 2005.

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

[4] S. Cameron. The economics of crime deterrence: A survey of theory and evidence. *Kyklos*, 41(2):301–323, 1988.

[5] T. Cymru. The underground economy: Priceless. *;login: The USENIX Magazine*, 31(6), 2006.

[6] D. E. Denning. Activism, hacktivism, and cyberterrorism. In J. Arquilla and D. Ronfeldt, editors, *Networks and Netwars*, pages 239–288. RAND, Santa Monica, CA, 2001.

[7] N. Farrel. Chinese hackers target pro-tibetan groups. *The Enquirer*, April 2008.

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

[9] J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents, July 2008. To appear.

[10] K. J. Higgens. *Dark Reading*, April 2007.

[11] G. Kessler. Defenses against distributed denial of service attacks. 2000.

[12] J. Leyden. Us credit card firm fights ddos attack. *The Register*, September 2004.

[13] D. Moore and C. Shannon. The caida dataset on the code-red worms. August 2001.

[14] R. Naraine. Kraken botnet infiltration triggers ethics debate. *eWeek.com*, May 2008.

[15] V. Paxson. An analysis of using reflectors for distributed denial-of-service attacks. *ACM Computer Communications Review (CCR)*, 31(3), July 2001.

[16] B. Potter. Dirty secrets of the security industry, 2007. Presented at Defcon XV (Las Vegas 2007).

[17] S. Schechter and M. Smith. How much security is enough to stop a thief? In *The Seventh International Financial Cryptography Conference (Gosier, Guadeloupe. January, 2003)*, 2003.

[18] StratFore. Situation report: Cyberwarfare and botnets. April 2008.

[19] I. Traynor. Russia accused of unleashing cyberwar to disable estonia. *The Guardian*, May 2007.

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

[21] N. Weaver and V. Paxson. A worst-case worm, 2004. To Appear in the Proceedings of The Third Annual Workshop on Economics and Information Security (WEIS04).

[22] N. Weinberg. Botnet economy runs wild. *Network World*, April 2008.

[23] C. C. Zou, W. Gong, and D. Towsley. Code red worm propagation modeling and analysis. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, pages 138–147, New York, NY, USA, 2002. ACM.

APPENDIX

A. RESULTS

H_e	One Shot Ceiling	Repeated Ceiling
1 player	b/L	b/L
$\frac{1}{N}\sum e$	Nb/L	b/L
min(e)	b/L	b/L
max(e)	None	0
kWT	None	None
kWTwM	b/L	b/L