

# Small World: Collisions Among Attackers in a Finite Population

Cormac Herley  
Microsoft Research  
One Microsoft Way  
Redmond, WA, USA  
cormac@microsoft.com

## ABSTRACT

The Internet is large, but it is finite. We examine the case of several attackers seeking victims in a large population where there is some prior statistic indicating likely viability. We show that collisions reduce attacker gains: the sum of what several attackers can extract is always less than what a single attacker would extract alone. The problem gets worse as the density of victims in the population decreases.

## 1. INTRODUCTION

The path from security vulnerabilities to money appears simple. An attacker picks a target from the Internet population, gains access to a resource, monetizes that access, and then repeats *ad infinitum*. Scale makes it appear plausible that this is easy and generates a lot of money. It is often suggested that this can all be automated and even tiny success rates can still produce great sums for criminals. For example, Menn says of phishing [11] “it didn’t take many for the math to work. Even if only one person in a hundred was a customer, millions would get the bait and several thousand of them would bite.” Schneier says [17]:

*“If you had a great scam to pick someone’s pocket, but it only worked once every hundred thousand tries you’d starve before you robbed anyone. In cyberspace, you can set your computer to look for the one-in-a-hundred-thousand chance. You’ll probably find a couple dozen a day. If you enlist other computers, you might get hundreds.”*

This sounds simple, but is perhaps a little too simple. A constant supply of new victims every day, and a lack of collisions with others seeking the same victims suggests that the population is unlimited and the attack never saturates. The same assumptions underpin multi-level marketing schemes. We show that this view is unrealistic.

Attackers naturally seek the best targets. They prefer to attack where they are most likely to succeed. A

natural model is that attackers judge how likely users are to be viable based on what they can observe. Everything they can learn about a potential victim might be wrapped into a single score,  $x$ , and their experience then allows them to estimate the probability of success,  $P\{\text{viable} \mid x\}$ . For example, the indicators of viability might include zip code, address, profession and anything that is known about likely wealth. Naturally, they would use this prior in deciding which users to attack. It makes a lot more sense to attack those most likely to be viable first. This is especially true for expensive attacks, where to be viable, a successfully attacked target has to pay handsomely.

A single attacker with the field to himself, can attack in descending order of likely viability: that is proceed from best to worst. He picks the most obvious targets and the goes after ones that are progressively less likely. At some point he can cease if things have decayed to the point where likely viability is too low for profitability, or if better opportunities present themselves elsewhere.

However, a single attacker seldom has the field to himself. Most attack opportunities are very competitive, with many using similar techniques and possessing similar information. Phishing isn’t the preserve of one attacker or gang. There are many involved in Nigerian 419-style scams, spam and malware distribution. Technique propagation guarantees that any new attack or scam gets to be used by many, not just the originators [17]. Thus, for most attacks, there are many who seek to profit from the pool of vulnerable victims.

We show that this radically changes the dynamics of the problem from the attackers’ perspective. Instead of attacking each target at most once (as the single attacker would do) there is now the risk of unintentional duplicated effort. Each attacker faces the risk that others (also seeking those most likely to be viable) will have been there first. This, of course, reduces everyone’s returns.

This is an example of a resource contention problem. Since attackers are independent, and there is no central control, the inevitable collisions reduce efficiency. Just

as Medium Access Control (MAC) protocols, such as ALOHA, seek to minimize collisions, we examine the economic best-case from the attackers viewpoint. We find that the economic value extracted by several independent actors is always less than the value a single attacker might extract. This factor difference can be as low as  $2\times$  when viable victims are plentiful, but rises sharply as density falls.

Press, examines the related question of screening for terrorists in a large population. There also, the sought individuals are rare, there also a memory-less sampling with replacement (which implies collisions) is used. Press [13] claims an interesting and counter-intuitive result, which is that selecting individuals to screen based on  $P\{\text{viable} | x\}$  does no better than random (in the limit of arbitrarily many screenings). Press proposes square-root sampling to minimize the number of searches per target found. Press' solution, however, ignores an economic constraint on attackers. When we constrain attackers to a minimum success rate we find that the difference between the obvious strategy and square-root sampling is not as great as in the resource-unconstrained case.

## 2. SAMPLING WITH REPLACEMENT

### 2.1 Victim distribution model

We consider a population of  $N$  users, which contains  $d \cdot N$  viable targets, where  $d$  is density. By viable we mean that these targets always yield a net profit when attacked, while non-viable targets yield nothing.

We assume that some users are far more likely to be viable than others. Viability is not directly observable: the attacker doesn't know with certainty that he will succeed unless he tries the attack. Nonetheless, the fact that some users are better prospects than others is observable. We assume that the attacker has a simple score,  $x$ , that he assigns to each user. The larger the score, the more likely in the attacker's estimate the user is to be viable.

More formally, the score,  $x$ , is a sufficient statistic [18]. The attacker might have several observations about the user, where he lives, his place of work, the accounts he possesses, *etc*: all of these be reduced to the single numeric quantity  $x$ . This encapsulates all of the *observable* information about the viability of user( $i$ ). Without loss of generality we'll assume that viable users tend to have higher  $x$  values than non-viable ones. This does not mean that all viable users have higher values than non-viable ones. For example, we might have  $pdf(x | \text{non-viable}) = \mathcal{N}(0, 1)$  and  $pdf(x | \text{viable}) = \mathcal{N}(\alpha, 1)$ . Thus, the observable  $x$  is normally distributed with unit variance, but the mean,  $\alpha$ , of  $x$  over viable users is higher than over non-viable users. An example is shown in Figure 1 (a).




Figure 1: (a) Distribution of  $x$  for normally distributed scores. The mean over non-viable users is zero (left-most curve). Various assumptions of the separation between viable and non-viable users are given. Means of  $\alpha = 1.18, 2.32$  and  $3.28$  are used. These result in the classifiers which have 90%, 95% and 99% ability to tell randomly chosen viable users from non-viable [7]. (b) The distribution of  $p_i = P\{\text{viable} | x_i\}$  for these three cases.

If viable and non-viable users are drawn from different distributions, as in Figure 1 (a), then from Baye's rule:

$$P\{\text{viable} | x\} = \frac{1}{1 + \frac{P\{x | \text{non-viable}\}}{P\{x | \text{viable}\}} \cdot \frac{1-d}{d}}$$

We define  $p_i \triangleq P\{\text{viable} | x_i\}$ , where  $x_i$  is the observed value of  $x$  for user( $i$ ). For convenience we'll assume that the  $p_i$  are sorted in descending order, so that for all values of  $i$  we have  $p_i \geq p_{i+1}$ . Figure 2 (b) shows  $p_i$  for the three different distributions shown in (a). The larger  $\alpha$  the greater the separation between  $P\{x | \text{non-viable}\}$  and  $P\{x | \text{viable}\}$  and the more skewed the distribution  $p_i$  becomes.

### 2.2 Single attacker

If an attacker has the field to himself there is no better strategy than to attack users in order of their likelihood to be viable. That is, he can observe  $x_i$  for each user( $i$ ), and then attack in decreasing order of  $p_i$ . The most likely victims are attacked first, and the probability drops as the population of victims is progressively exhausted. Thus:

$$P\{\text{viable on attack-}k\} = p_k.$$

The average number of attacks per viable victim found is:

$$\mu_{solo} = \sum_{i=0}^{N-1} p_i \cdot i. \quad (1)$$

### 2.3 Several attackers

It is unrealistic to assume that anyone gets to reserve an opportunity to himself. Now what happens when there is not a single attacker, but several?

Suppose that we have several attackers. We assume that once a viable victim is attacked he ceases to be viable, but this fact is not observable. That is, the features that made him an attractive target in the first place (*e.g.*, address, profession, *etc*) are unaltered even though his stealable money is now gone. Thus,  $x_i$  remains unchanged after a successful attack. For example, if  $p_i$  is large, then user( $i$ ) is a tempting target; but, if another attacker has been there first, user( $i$ ) is no longer viable.

We assume that attackers are independent, so they have no opportunity to co-ordinate their attacks or agree who will attack where. We assume that all of the attackers have the same information and abilities. We are interested in the best case for what attackers as a group can achieve, thus our results will provide an upper-bound on their impact. Thus, our goal is not to find the strategy that gives one attacker advantages over his peers: any strategy available to one is available to all. However, even though they compete, the attackers have a common interest in reducing the number of wasted attacks. We will revisit these assumptions in Section 4.2.

We address the question of efficiency in the face of resource contention. Several resource contention strategies that are common in other domains seem inapplicable here. For example, locking strategies and carrier sensing (techniques used in Medium Access Control) require protocols agreed-upon in advance, which seems inappropriate for attackers who are not co-operating [10]. Collision detection (another common strategy) doesn't help if the full cost of a collision (*i.e.*, attack) has to be incurred before detection is possible. The common technique that seems appropriate is randomization of access (such as is used in MAC protocols). We seek a strategy to minimize collisions so that each attacker independently selects targets, without knowing where others

have attacked. Thus, the problem is one of sampling-with-replacement.

We'll denote by  $q_i$  the probability that user( $i$ ) is attacked on the  $i$ -th attack (by any attacker). (Note: The problem is not precisely one of sampling with replacement: each attacker knows where he has attacked before, but not where his peers have. We will assume that there are enough attackers so that this effect is minor. It can be shown that is approximately true, even for five to ten attackers.) The probability that user( $i$ ) is attacked on the  $k$ -th attack, has not been attacked before and is viable is:

$$q_i(1 - q_i)^{k-1}p_i.$$

Summing over the whole population then gives the probability that a viable user is found on the  $k$ -th attack:

$$P\{\text{viable on attack-}k\} = \sum_{i=0}^{N-1} q_i \cdot (1 - q_i)^{k-1}p_i. \quad (2)$$

The average number of attacks per found viable user is (if we allow the attacks to continue to infinity):

$$\begin{aligned} \mu &= \sum_{k=0}^{\infty} k \cdot \sum_{i=0}^{N-1} q_i \cdot (1 - q_i)^{k-1}p_i \\ &= \sum_{i=0}^{N-1} q_i \cdot p_i \sum_{k=0}^{\infty} k \cdot (1 - q_i)^{k-1} \\ &= \sum_{i=0}^{N-1} \frac{p_i}{q_i}. \end{aligned} \quad (3)$$

### 2.4 Sampling strategies

We now examine several strategies for choosing the  $q_i$ .

#### 2.4.1 Uniform sampling

In uniform sampling attackers simply choose targets independently of the probability that the user is viable. That is  $q_i = 1/N$ . This certainly reduces the collisions between attackers, as nobody is concentrating their attention on the best targets, and everyone addresses the whole population equally. In the case of uniform sampling (with replacement), the average number of attacks per viable victim found is (from (3)):

$$\mu_{unif.} = \sum_{i=0}^{N-1} \frac{p_i}{1/N} = N.$$

#### 2.4.2 Importance sampling

Attacking uniformly at random seems likely to be sub-optimal. If we have any information about viability, it seems better to use it. A very obvious alternative approach is to attack user( $i$ ) with probability  $p_i$ . That is,  $q_i = p_i$ , attack a user with probability proportional to the *a priori* probability that he is viable. Surprisingly,

	$\alpha = 1.18$	$\alpha = 2.32$	$\alpha = 3.28$
$\mu_{solo}$	49,375	15,136	5,133
$\mu_{unif.}$	$10^6$	$10^6$	$10^6$
$\mu_{import.}$	$10^6$	$10^6$	$10^6$
$\mu_{sqrt}$	261,353	96,985	16,085

**Table 1: Improvement of square-root sampling with respect to importance, or uniform. The table shows the average number of attacks required for various strategies. These values are for a population of  $N = 10^6$  and viable victim density of  $d = 0.01$ , and the three distributions shown in Figure 1 (a). Note that the improvement that both a solo attacker, and square-root sampling realize is a function of concentration: when  $\alpha$  is large, and  $p_i$  is very skewed, some targets are very obvious.**

as observed originally by Press [13], the average number of attacks becomes (so long as none of the  $q_i$  are zero):

$$\mu_{import.} = \sum_{i=0}^{N-1} \frac{p_i}{q_i} = \sum_{i=0}^{N-1} \frac{p_i}{p_i} = N. \quad (4)$$

Thus, the average number of attacks under importance sampling (with replacement) is exactly the same as attacking uniformly at random. This might appear puzzling: is the information  $p_i = P\{\text{viable} \mid x_i\}$  of no help? It hardly seems right that ignoring the prior information about viability can be best.

### 2.4.3 Square-root sampling

To overcome the failure of importance sampling to do better than uniform Press [13] suggests choosing the  $q_i$ 's to minimize the average number of attacks per victim found. That is minimize (3) subject to the constraint  $\sum_i q_i = 1$ . This turns out to be:

$$q_j = \frac{\sqrt{p_j}}{\sum_{i=0}^{N-1} \sqrt{p_i}}. \quad (5)$$

We will refer to this choice as square-root sampling. Using this choice, we find the mean number of attacks per victim found is:

$$\mu_{sqrt} = \left( \sum_{i=0}^{N-1} \sqrt{p_i} \right)^2. \quad (6)$$

This can be a considerable improvement, as illustrated in Table 1.

## 2.5 Convergence to uniform

Figure 2 shows the probability of finding a viable victim as the number of attacks grows for various strategies. That is, it shows  $P\{\text{viable on attack-}k\}$  vs.  $k$  for the different choices of  $q_i$ . The reciprocal of the probability would be the expected number of attacks per

viable victim at a certain point. This is graphed for solo, uniform, importance, and square-root sampling. Clearly, importance and square-root do very well at first, but performance degrades. At some point they do even worse than uniform: as (4) shows, what importance sampling gains in the early stages it loses later on, doing no better than uniform sampling in the limit.

Press claims the advantage of square-root sampling over importance sampling as support for the claim that relying too heavily on the prior is a mistake. This results in the same ‘‘obvious’’ targets being attacked over and over, with great loss of efficiency.

The problem with importance sampling becomes evident when examining Figure 2. Importance sampling does much better at first, but once the easy targets have been found it actually does worse. There are some users who have very low values  $p_i$  and yet are viable. Importance sampling, where  $q_i = p_i$ , has great difficulty finding them. This is similar to the Coupon Collector’s problem: the first coupons’s may be found quickly, but sampling-with-replacement becomes more and more wasteful as we chase the last few coupons.

Square-root sampling, as suggested by Press, improves things somewhat, but also does worse than uniform sampling eventually. However, the problem is only manifest when we seek to extract all victims, and use the average number of attacks as criterion. If we are willing to halt earlier, we may be able to do a great deal better.

## 3. ECONOMICALLY MOTIVATED ATTACKERS

A weakness of the approaches above, is that we assumed attacks persist to infinity. This led to the simple form (3) for attacks per target found; and choosing the  $q_i$  to minimize (3) led to square-root sampling. While this might be appropriate to Press’ goal of seeking terrorists it is clearly not appropriate when we consider an economically motivated attacker. It makes sense to persist to infinity only if we are determined to find every viable victim and have no resource constraint. It seems more realistic, however, to assume that attackers will persist only so long as it is profitable.

Consider the evolution of the probability of finding a viable victim shown in Figure 2. As the number of attacks,  $k$ , increases each of the solo, importance and square-root strategies decay in quality. In fact, each of them eventually does worse than uniform. It seems very unlikely that any attacker will persist when the likelihood of finding a victim falls by 4 (importance and square-root) or 9 (solo) orders of magnitude from what he began with. If we assume that attackers are economically motivated it makes sense that they will have a threshold minimum probability of success,  $p_{min}$ . For example, if the average gain divided by the average cost were 20 then the probability of finding a viable victim

should be no lower than  $1/20 = 0.05$  on average. Thus, it makes little sense to persist with any strategy after  $P\{\text{viable on attack } k\} < p_{min}$ .

This suggests that choosing the  $q_i$  to minimize (3), as square-root sampling does, is sub-optimal. When  $p_i$  is low enough that  $P\{\text{viable on attack } k\} < p_{min}$  there may be some viable users, but on average attackers will lose more than they gain in pursuing them. It makes more sense to exclude these from consideration and optimize only over the portion of the population that at least has a chance of being profitable.

Thus, we assume that once the probability of success drops below  $p_{min}$  attacks stop. In the case of the solo attacker this means:

$$\mu'_{solo} = \sum_{i=0}^{i_{max}} p_i \cdot i,$$


where  $i_{max}$  is the largest  $i$  such that  $p_i \geq p_{min}$ .

Equally, we define  $k_{max}$  as the largest  $k$  such that  $P\{\text{viable on attack-}k\} \geq p_{min}$ . Then the average number of attacks per viable user found becomes:

$$\begin{aligned} \mu' &= \sum_{k=0}^{k_{max}} k \cdot \sum_{i=0}^{N-1} q_i \cdot (1 - q_i)^{k-1} p_i \\ &= \sum_{i=0}^{N-1} q_i \cdot p_i \sum_{k=0}^{k_{max}} k \cdot (1 - q_i)^{k-1} \\ &= \sum_{i=0}^{N-1} \frac{p_i}{q_i} \cdot (1 - (1 - q_i)^{k_{max}}). \end{aligned} \quad (7)$$

We now examine how the various strategies do under this measure of success. We define inefficiency as  $\mu' / \mu'_{solo}$ , the ratio of average number of attacks per viable victim found to the average for the solo attacker. We graph this for importance and square-root sampling in Figure 3 as a function of density for the least skewed (*i.e.*,  $\alpha = 1.18$ ) distribution (the other distributions show a similar trend). Clearly, when victims are very plentiful the inefficiency is low: competing attackers harvest victims at an efficiency that is only a factor of 2-3 worse than the solo attacker does. However, things get rapidly worse: at a victim density of  $d = 10^{-4}$  both importance and square-root strategies are doing about a factor of 100 worse.

In Figure 4 we show the cumulative fraction of victims found (*i.e.*, number found divided by total) as the number of attacks increases. Again  $N = 10^5$  is used in a Monte Carlo simulation in which the least skewed distribution from Figure 1 (a) was used. Figure 4 (a) uses a density of  $d = 10^{-2}$ , (b) uses  $d = 10^{-3}$  and (c) uses  $d = 10^{-4}$ . Observe that the competing attackers (importance (blue) and square-root (red)) always do worse than the solo attacker (black): they always find fewer victims per attack. For example, Figure 4 (a) shows that at  $d = 10^{-2}$  to find 40% of victims competing at-




**Figure 2: Probability of finding a viable victim (*i.e.*, equation (2)) as number of attacks,  $k$ , increases. The least skewed of the distributions from Figure 1 (a) is used (*i.e.*,  $\alpha = 1.18$ ). Four strategies are shown: solo (black), importance (blue), square-root (red) and uniform-at-random (green) sampling. This involves  $N = 10^5$  users and  $d = 10^{-3}$ . Observe that all of the strategies eventually become worse than random.**

tackers using square-root sampling need  $> 10\times$  more attacks per success than the solo attacker. The penalty with respect to the solo attacker gets worse as the density of victims decreases: Figure 4 (c) shows a widening margin between the strategies.

Figure 4 (c) also makes clear the disparity in the number of victims found for a fixed attack budget. At this density, with a total attack budget of  $\log_{10}(k/N) = -3$  (*i.e.*, attacking one in a thousand) the solo attacker finds 25% of victims while the importance and square-root strategies for competing attackers find 9% and 3% respectively.

Observe that the gap between the solo and competing strategies is worst when  $k/N$  is small. A very plausible strategy for the solo attacker is to pick only the most likely targets and then stop. The cumulative fraction of victims that he will find this way is small, but he can do so at very low cost. Figure 4 (c) makes clear that competition is fiercest for the earliest, and most obvious, victims. The solo attacker can get the easiest 10% of victims for about  $100\times$  less effort than the square-root strategy.

Finally, observe that Figure 4 provides little support for the claim that a square-root strategy is better than importance sampling. While it may be better in the limiting case where we pursue every victim, it appears worse than importance sampling at most operating points of interest. It avoids the waste in the late stages of a coupon collector's problem, but appears to do worse early on.



**Figure 3: Inefficiency,  $\mu' / \mu'_{solo}$  as a function of density. Exploitation of an opportunity by several independent attackers is always less efficient than when exploited by one. Monte Carlo simulation with  $N = 10^6$  users and attackers who need an average success rate of  $p_{min} \geq d$ .**

## 4. DISCUSSION


### 4.1 Comparison to [7]

We find that several attackers always make less efficient use of an opportunity than a solo attacker would. Further, the inefficiency increases rapidly as victim density falls. The latter conclusion echoes our earlier result that economic value falls far faster than victim density [7]. It is worth emphasizing that these are not alternative paths to the same conclusion. The result of [7] was that the economic opportunity for a solo attacker dropped very sharply with  $d$ . We have here shown that several attackers extract less value than one, and again the situation gets worse with  $d$ . It appears the deterioration for multiple attackers is then catastrophic as density falls: they can extract only a falling fraction of a falling total. This reinforces the conclusion reached earlier, when attacks with low enough victim densities pose little economic threat.

### 4.2 Equal attackers

Our model assumed attackers who were equal in that they had the same information and abilities. None had the ability to increase his share at the expense of others. Thus, the best they could do (in the absence of central control or collusion) is minimize the overall waste. Under these assumptions, as we have seen, waste is inevitable. No matter how good the original opportunity things converge quickly to the returns that would be found attacking at random.

Under these assumptions we have shown that the economic opportunity decays as others join the party. However, it would be a mistake to think that attacks



**Figure 4: Cumulative fraction of victims founds as the number of attacks increases. Monte Carlo simulation with  $N = 10^5$ ; the least skewed distribution from Figure 1 (a) is used. The results for solo (black), importance (blue) and square-root (red) sampling are shown. Observe that several attacker always do worse than a solo attacker: more total attacks are needed to find the same fraction of victims. Note that the gap between the performance of the solo attacker and the colliding attackers is a function of victim density: the lower the density the worse the effect of collisions on return. (a) Density  $d = 10^{-2}$  (b) Density  $d = 10^{-3}$  (c) Density  $d = 10^{-4}$ .**

on small victim densities do not occur. Several attackers with the same skills and knowledge rapidly destroy their communal opportunity. Things are obviously a great deal better for those who have unique skills, or unique knowledge about viability. Those who have only commodity skills and knowledge share only bad opportunities.

## 5. RELATED WORK

The question of attacker collisions has occasionally been raised, but has seldom been explored in detail. Geer and Conway allude to it [3]. Enright *et al.* [2] mention the possibility of different research teams colliding while examining the same botnet. The problem does not appear to have received a systematic treatment.

Herley examines the difficulty of attackers who target as a strategy. He shows that competing against a scalable attacker is hard [6]. He shows that economic opportunity drops much faster than victim density [7]. Thus, an opportunity with victim density  $d/2$  has less than half the value of one with density  $d$ . This implies that opportunities with very low densities are hard to exploit.

A point of contrast between this work and [7] is that there a single attacker makes a binary decision as to whether to attack or not. Here, several attackers compete. The main finding of [7] was that the economic opportunity falls far faster than density. In this work we find that the inefficiency that several attackers experience increases with density.

Press examines a problem very close to ours, where rare malfactors are sought in a large population. The sampling is constrained to be done with replacement. Press suggests square-root sampling as the way of minimizing the expected number of searches before malfactors are found. This criterion however, assumes that screening can continue to infinity (*i.e.*, there's no resource constraint).

Variations of the Coupon Collector's problem occur in numerous engineering problems. The last-block problem in coding. Maximizing the use of a resource under collisions occurs in many areas. Medium Access Control (MAC) policies such as ALOHA try to maximize throughput subject to the constraint that an unknown number of users compete for the resource.

Herley and Florêncio [8] examine the economic competition that leads to a Tragedy of the Commons.

Anderson [14] shows that incentives greatly influence security outcomes and demonstrates some of the perverse outcomes when they are mis-aligned. Since 2000 the Workshop on the Economics of Information Security (WEIS) has focussed on incentives and economic tradeoffs in security.

Varian suggests that many systems are structured so that overall security depends on the weakest-link [5].

Gordon and Loeb [9] describe a deferred investment approach to security. They suggest that, owing to the defender's uncertainty over which attacks are most cost effective, it makes sense to "wait and see" before committing to investment decisions. Boehme and Moore [15] develop this approach and examine an adaptive model of security investment, where a defender invests most in the attack with the least expected cost. Interestingly, in an iterative framework, where there are multiple rounds, they find that security under-investment can be rational until threats are realized. Unlike much of the weakest-link work, our analysis focusses on the attacker's difficulty in selecting profitable targets rather than the defender's difficulty in making investments. However, strategies that suggest that under-investment is not punished as severely as one might think spring also from our findings.

Grossklags *et al.*[4] examine security from a game theoretic framework. They examine weakest-link, best-shot and sum-of-effort games and examine Nash equilibria and social optima for different classes of attacks and defense. They also introduce a weakest-target game "where the attacker will always be able to compromise the entity (or entities) with the lowest protection level, but will leave other entities unharmed." A main point of contrast between our model and the weakest-target game is that in our model those with the lowest protection level get a free-ride. So long as there are not enough of the to make the overall attack profitable, then even the weakest targets escape.

Fultz and Grossklags [12] extend this work by now making the attacker a strategic economic actor, and extending to multiple attackers. As with Grossklags *et al.*[4] and Schechter and Smith [16] attacker cost is not included in the model, and the attacker is limited mostly by a probability of being caught. Our model, by contrast, assumes that for Internet attackers the risk of apprehension is negligible, while the costs are the main limitation on attacks.

Odlyzko [1] addresses the question of achieving security with insecure systems, and also confront the paradox that "there simply have not been any big cybersecurity disasters, in spite of all the dire warnings." His observation that attacks thrive in cyberspace because they are "less expensive, much more widespread, and faster" is similar to our segmentation of broadcast attacks.

## 6. CONCLUSION

We examine how independent attackers share an economic opportunity. The collisions that inevitably arise reduce returns, so that the value extracted is always less than a solo attacker would do. This factor gets significantly worse as victim density falls. Competition is worst for the easiest targets, *i.e.*, those most likely to

be viable. Thus, concentrating on only the easiest targets (a very plausible strategy when there is only one attacker) is the very point where competition is fiercest when there are many.

## 7. REFERENCES

- [1] A. Odlyzko. Providing Security With Insecure Systems. *WiSec*, 2010.
- [2] B. Enright, G. Voelker, S. Savage, C. Kanich, and K. Levhchenko. Storm: when researchers collide. *;login*, 2008.
- [3] D. E. Geer Jr and D. G. Conway. Beware the ids of march. *IEEE Security & Privacy*, page 87, 2008.
- [4] J. Grossklags, N. Christin, and J. Chuang. Secure or insure?: a game-theoretic analysis of information security games. *WWW*, 2008.
- [5] H. R. Varian. System Reliability and Free Riding. *Economics of Information Security*, 2004.
- [6] C. Herley. The Plight of the Targeted Attacker in a World of Scale. *WEIS 2010, Boston*.
- [7] C. Herley. Why do Nigerian Scammers say they are from Nigeria? *WEIS 2012, Berlin*.
- [8] C. Herley and D. Florêncio. A Profitless Endeavor: Phishing as Tragedy of the Commons. *NSPW 2008, Lake Tahoe, CA*.
- [9] L.A. Gordon and M.P. Loeb. The Economics of Information Security Investment. *ACM Trans. on Information and System Security*, 2002.
- [10] E. A. Lee and D. G. Messerschmitt. Digital communications. *Kluwer Academic Publishers*, 1994.
- [11] J. Menn. *Fatal System Error: The Hunt for the New Crime Lords Who Are Bringing Down the Internet*. PublicAffairs, 2010.
- [12] N. Fultz and J. Grossklags. Blue versus Red: Toward a Model of Distributed Security Attacks. *Financial Crypto*, 2009.
- [13] W. Press. Strong profiling is not mathematically optimal for discovering rare malfeasors. *Proceedings of the National Academy of Sciences*, 106(6):1716, 2009.
- [14] R. Anderson. Why Information Security is Hard. In *Proc. ACSAC*, 2001.
- [15] R. Boehme and T. Moore. The Iterated weakest-link: A Model of Adaptive Security Investment. *WEIS*, 2009.
- [16] S. Schechter and M. Smith. How Much Security is Enough to Stop a Thief? In *Financial Cryptography*, pages 122–137. Springer, 2003.
- [17] B. Schneier. *Secrets and lies: digital security in a networked world*. Wiley, 2000.
- [18] H. L. van Trees. *Detection, Estimation and Modulation Theory: Part I*. Wiley, 1968.