

A Profitless Endeavor: Phishing as Tragedy of the Commons

Cormac Herley and Dinei Florêncio
Microsoft Research
One Microsoft Way
Redmond, WA, USA
c.herley@ieee.org, dinei@microsoft.com

ABSTRACT

Conventional wisdom is that phishing represents easy money. In this paper we examine the economics that underly the phenomenon, and find a very different picture. Phishing is a classic example of tragedy of the commons, where there is open access to a resource that has limited ability to regenerate. Since each phisher independently seeks to maximize his return, the resource is over-grazed and yields far less than it is capable of. The situation stabilizes only when the average phisher is making only as much as he gives up in opportunity cost.

Since the picture we paint is at variance with accepted wisdom we check against several publicly available data sources on phishing. We find the oft-quoted survey-based estimates of phishing losses unreliable. In particular the victimization rate found in most surveys is smaller than the margin of error, and dollar losses are estimated by averaging unverified self-reported numbers. We estimate that recent public estimates overstate phishing losses by as much as a factor of fifty.

This economic portrait illuminates our enemy in an entirely new light. Far from being a path to riches, phishing appears to be a low-skill low-reward business. The enormous amount of phishing activity is evidence of its failure to deliver riches rather than its success, as phishers send more and more email hoping for their share of the bounty that eludes them. Repetition of questionable survey results and unsubstantiated anecdotes makes things worse by ensuring a steady supply of new entrants.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSPW'08, September 22–25, 2008, Lake Tahoe, California, USA.
Copyright 2008 ACM 978-1-60558-341-9/08/09 ...\$5.00.

Categories and Subject Descriptors

K.6.0 [Economics]:

General Terms

Economics

Keywords

Phishing, fraud, identity theft

1. INTRODUCTION

Phishing has pushed its way to the forefront of the plagues that confront online users. While accurate studies have documented the enormous growth in the amount of phishing email sent and number of phishing sites reported [22], estimates of the amount that phishers make have been harder to come by. This is not astonishing, as illegal enterprises file no taxes, and do not report to the USA Securities and Exchange Commission (SEC). Thus we do not have good answers to the most fundamental “Know Your Enemy” questions. How much does a phisher make, and what’s the total amount of money stolen per year? At first glance, phishing looks like a very profitable business for an individual. There is little capital outlay or startup costs, no raw materials and no sophisticated equipment to rent or buy. The phisher merely harvests “free money” from the online population. This would appear to compare very favorably with other businesses that involve significant investment or skill or training. Accounts in the popular and web press support this view. Reports of the ease with which money can be made tend to be sensational. An interview with a Phisher [1], for example, tells of an 18 year old who claims to have stolen “way over 20 million identities” and to make \$3-4k per day. The NY Times (July 4, 2006) had a headline “Identity thief finds easy money hard to resist.” Thomas and Martin [23] tell a compelling story about life in the underground economy and claim “even those without great skills can barter their way into large quantities of money they would never earn in the physical world.” They describe

sophisticated divisions of labor, alliances and techniques for building trust, and quote numerous snippets of negotiations on IRC networks. While all of the accounts are anecdotal, and no verification is offered, the reader of [23] is left with the impression that phishing, and the underground economy in general, represents very easy money.

Yet there is something very wrong with this picture: common sense dictates that low-skill jobs pay like low-skill jobs, whether the activity is legal or not. Phishing requires basic computer skills; that should pay better than minimum wage, but there is no large barrier to entry (the phisher in the NY Times story bought the software he needed for \$60). Do phishers live by a set of economic laws different from those the rest of us experience? If phishers made as much as surgeons wouldn't new entrants increase competition and drive the returns down? Suppose there were a fixed number of dollars available to be phished each year; that fixed pool would be divided among more and more people and each phisher's take would decrease. New entrants stop arriving only when the opportunity is no better than the opportunities elsewhere. So this argues that a fixed pool would be divided among a community of phishers that expands to drive the average return down. So far so obvious.

However, as we will show, the economics of phishing are far far worse than this. Rather than sharing a *fixed* pool of dollars phishing is subject to the tragedy of the commons [16]; *i.e. the pool of dollars shrinks as a result of the efforts of the phishers*. A community (all phishers) share a finite resource (the pool of phishable dollars) that has limited ability to regenerate (dollars once phished are not available to other phishers). The tragedy of the commons is that the rational course of action for each individual (phisher) leads to over-exploitation and degradation of the resource (the phishable dollars).

So what are the economics of this situation, where a resource that has limited ability to regenerate is accessible by anyone who wishes to exploit it? This question it turns out has been asked and answered very thoroughly in the Economics literature. In a classic paper Gordon [15] examined the economics of a Common Property resource. The example treated by Gordon is that of the fishing industry (*i.e.* "real" fishing where men go to sea on trawlers and return with cod, haddock *etc*). It emerges that whoever named phishing chose well as the economics of fishing and phishing have a great deal in common. In both cases there is a predator-prey relationship: between fishermen and cod, or between phishers and dollars. In both cases the prey has a limited ability to regenerate (neither fish nor dollars have infinite capacity for growth). In both cases there is *open access* to the prey: neither fishermen nor

phishers have the opportunity to erect a fence and restrict access. This leads to the tragic overgrazing of the commons: the resource yields far less when exploited by independent actors than if it were managed by a single decision maker. Of course this is "tragic" only for the phishers, the degradation of their commons means that dollar losses due to phishing fall. The over-exploitation predicted by theory has been verified in several open access markets. This is the case in fishing [15, 25], piracy and privateering [7], mugging [20] and grassland exploitation [2].

In this paper we will apply basic open access economic theory to phishing. The picture that we end up with is very different from the "easy money" that is conventional wisdom. A brief summary is that the total revenue (all dollars stolen through phishing) is equal to the total cost (dollar value of the opportunity that phishers gave up in other occupations). The average revenue for a given phisher is the same (or slightly lower) than he would have made at another available occupation for his skill level. The easier phishing gets the worse the economic picture for phishers. As phishers put more and more effort into the endeavor the total revenue falls rather than rises. This last point is particularly interesting as it suggests that the increasing volumes of effort measured in [22] indicate decreasing rather than increasing total revenue. The fact that PayPal's CSO stated in Feb. 2007 [3] that phishing "is not even in the top five" loss threats that PayPal faces leads us to believe that our analysis has merit.

The next section reviews the economics of open access resources and applies the model to phishing. Section 3 looks at the implications that this has in determining the enemy we face and how to fight them. In Section 4 we examine several objections that might be raised to this analysis. Section 5 reviews related work, and in particular examines the publicly available sources of data on the dollar size of the phishing problem. Section 6 concludes.

2. THE ECONOMICS OF OPEN ACCESS RESOURCE POOLS

We will closely mimic the notation used by economists modelling open access fishing grounds. See for example [15, 4].

2.1 Sustainable Harvest as a Function of Sustained Effort

This Section shows how Figure 1, which relates the sustainable harvest with the sustained effort, is derived. Those familiar with Economics, or willing to accept the common sense explanation of the figure can skip to Section 2.2. Let X be the total pool of phishable dollars. This is probably less than all dollars in all online accounts: some users are too savvy to fall for phish-

ing, some institutions may have extremely tight restrictions on wiring money from accounts. Much as Gordon [15] makes no attempt to estimate the number or total weight of fish in a fishing ground, we won't attempt to place a numeric value on X ; what matters is that it is finite. The analysis seeks to reveal the economic factors that cause equilibrium to be reached rather than estimate quantities. Let E be the total effort of all phishers; if the main resource a phisher has is his time we can measure E in hours. Let $H(X, E)$ be the total dollar harvest per unit time (which depends on the pool of available dollars and the total effort).

Unchecked the pool of phishable dollars grows over time. The growth is dependent on X itself. That is $\frac{dX}{dt} = f(X)$. This is the expected behavior of any quantity that has an exponential growth pattern, but is contained in a bounded resource pool. Thus $\frac{dX}{dt}$ grows as depicted in Figure 2 (a). The number of dollars added per unit time gets larger as X gets larger, but the growth slows and finally drops to zero when X has reached the resource limit and no further growth is possible. For example, as X approaches the total number of dollars in all online accounts $\frac{dX}{dt}$ must approach zero.

But, of course, our pool of phishable dollars does not grow unchecked: every dollar that is harvested by a phisher is removed from the pool. Thus the true growth rate of the pool is the unchecked rate minus the harvest:

$$\frac{dX}{dt} = f(X) - m \cdot H(X, E).$$

Actually, the phishable dollars are reduced *at least* by $H(X, E)$. Dollars stolen are removed from the pool, however, there is a possibility that each dollar stolen causes more than one dollar to leave the pool. This is so since a victim who has his PayPal credentials stolen (and loses money as a consequence) is likely to be especially careful with any remaining money (*e.g.* change the password if the account has not yet been emptied) and to be more alert with respect to any other accounts. We account for this factor by removing $m \cdot H(X, E)$ rather than $H(X, E)$ from the pool; clearly $m \geq 1$. Again, the actual value of m will not affect the analysis much, we can assume $m = 1$ if we choose.

The phishers can only sustainably harvest the growth rate at any X . That is, in equilibrium $\frac{dX}{dt} = 0$ and hence

$$f(X) = m \cdot H(X, E).$$

If $m \cdot H(X, E) > f(X)$ then the pool of phishable dollars shrinks to zero (*i.e.* phishers consistently harvest more than the replacement rate of the dollars). If $m \cdot H(X, E) < f(X)$ the dollar pool grows; however, as X increases at some point $f(X)$ begins to fall (as depicted in Figure 2 (a)). So if $m \cdot H(X, E) < f(X)$ then X will increase until $m \cdot H(X, E) = f(X)$. Thus, in equilibrium we have $\frac{dX}{dt} = 0$, which implies that, for

a given X , the sustainable harvest is

$$H(X, E) = \frac{1}{m} \cdot f(X).$$

At this sustained level of harvesting the pool neither increases nor decreases.

Now, in equilibrium, the pool of phishable dollars X depends on the effort E . When there is no phishing effort ($E = 0$) X achieves its maximum. For some large enough effort we will have $X = 0$ (*e.g.* if everyone is phished every day the pool of phishable dollars will be zero). In between those extremes X is inversely related to E . Following [15] we depict this as a linear relationship in Figure 2 (b), but it doesn't significantly change the analysis if it deviates from this. What matters is that the number of phishable dollars is a function of effort (*i.e.* $X = X(E)$) and decreases as E increases (Note in Figure 2 (b) $X(E)$ is the dependent variable).

Now the harvest that phishers extract from the pool is a function of effort E and the size of the pool: $H = H(X, E)$. But in steady state, as we have seen, X can be expressed as a function of effort. So $H(X, E) = H(X(E), E) = H(E)$ and the sustainable harvest can be expressed as a function of effort alone. This is done in Figure 2 (c). For example, at effort E_0 we can determine the sustainable harvest $H(E_0)$ by finding the phishable dollars for that level of effort $X_0 = X(E_0)$ and then equating $H(E_0) = 1/m \cdot f(X_0)$ (since we know that the harvest and the growth must be equal in equilibrium). Thus we end up with a curve that shows the sustainable harvest for any particular level of effort in Figure 2 (a). For convenience, this is reproduced in Figure 1.

2.2 Summary so far: Sustainable Harvest at Sustained Effort

While, there was some analysis involved in its derivation Figure 1 represents what common sense suggests. The sustainable harvest depends on the phishing effort. When $E = 0$ the harvest is also zero. As sustained effort increases so does the sustainable harvest. However, at some level of phishing effort, the sustainable harvest peaks and returns to zero. This must be so since, at some level of harvesting effort (*e.g.* everyone is phished every day), the pool of phishable dollars drops to zero (and hence so also must the harvest).

The curves used to derive this graph are for example only. The key assumptions are that the growth as a function of the pool of phishable (*i.e.* dX/dt vs. X) as in Figure 2 (a) falls to zero for some large enough X . And that the pool of phishable dollars against effort (*i.e.* $X(E)$ vs. E) as in Figure 2 (b) is monotonically decreasing.

2.3 Independent Profit Maximizing Actors

Figure 1 shows the sustainable harvest achievable as a function of effort. Any point on the curve is achievable




Figure 1: Open access exploitation of a finite resource. The curve $H(E)$ is the sustainable harvest at any level of sustained effort; the harvest represents the total revenue $R_{tot}(E)$. The total cost is proportional to the effort $C_{tot}(E) = a \cdot E$. Economic equilibrium is reached when $C_{tot}(E) = R_{tot}(E)$. Even though the resource could yield greater revenue, $H(E_x)$, with less effort, E_x , the tragedy of the commons causes overgrazing and destruction of the resource.

for some E . Now the total revenue per unit time (*i.e.* amount stolen by all phishers) is just equal to the harvest $R_{tot}(E) = H(E)$. The total cost is $C_{tot}(E) = a \cdot E$, where a comprises all costs per unit of effort *including opportunity costs* of the phishers. For example, if we measure E in hours, a would include the average amount that the phishers could make from an alternative use of the effort, plus the hourly cost of running their equipment, plus the depreciated value of any one-time investments they made. For most phishers the main component of this cost is probably the lost opportunity of other employment (*i.e.* a is how much they would have made per hour doing something else).

Clearly, if a single decision maker decided the total amount of effort, he would choose the value of E that maximizes his total profit: $H(E) - a \cdot E$. This occurs at or below the effort E_x at which $H(E)$ is maximum. For a single decision maker increasing the effort beyond E_x makes no sense: not only does his effort, and hence cost, go up, but his total revenue goes down (*i.e.* not merely revenue per unit of effort). Thus increasing effort above E_x is very self-destructive to the profit-maximizing interest. This, however, is where the tragedy of the commons enters the picture: no phisher gets to fence off the pool of phishable dollars and harvest it for himself. Nobody is in a position to limit the overall effort put into harvesting. Decisions are made by many independent actors each seeking to maximize their return.

Let us examine how the resource fares under this ar-

range. The average return that each phisher gets per unit of effort is $R_{avg}(E) = H(E)/E$. His cost for that unit of effort is a . The phisher makes a profit when $R_{avg}(E) > a$, a loss when $R_{avg}(E) < a$, and “breaks even” when $R_{avg}(E) = a$. That is, $R_{avg}(E)$ must exceed the money per unit of effort that the phisher could make in other employment for this to be profitable.

But this fails to happen. For example, suppose at a given overall effort E we have $R_{tot}(E) > C_{tot}(E)$. This gives $R_{avg}(E) > a$, and phishing is profitable for all of the participants: they make more than the opportunities they gave up elsewhere. But, since the opportunity is profitable, each phisher has the incentive to exert as much effort as possible. In addition, the profitability of the opportunity attracts new entrants to the pool of phishers. Thus the overall effort increases as each independent phisher seeks to maximize his return. However, from Figure 1 we see that increasing effort decreases $R_{tot}(E)$. The average return $R_{avg}(E) = H(E)/E$ drops even more sharply, as the numerator is decreasing while the denominator is increasing. Economic equilibrium is reached when there is no incentive to increase effort further [15, 4, 16]. This happens when $R_{avg}(E) = a$. This means that new entrants are not attracted, since phishing pays them no more than the opportunity they give up elsewhere, and existing phishers have no incentive to increase effort.

3. IMPLICATIONS FOR PHISHING

There are a number of interesting consequences that spring from the observation that phishing involves a tragedy of the commons.

First, even though it harvests “free money,” in economic equilibrium phishing generates total revenue equal to the total costs incurred by the actors. That is $R_{tot}(E) = C_{tot}(E)$. In addition, for the individual actors $R_{avg}(E) = a$: each participant earns, on average, only as much as he would have made in the opportunities he gave up elsewhere.

Second, *as the total phishing effort increases the total phishing revenue declines*. This leads to the most profound mis-allocation of effort: the harder individual phishers try the worse their collective situation gets. To quote Gordon [15]: “This is why fishermen are not wealthy, despite the fact that the fishery resources of the sea are the richest and most indestructible available to man. By and large, the only fisherman who makes it rich is the one who makes a lucky catch [...].”

Third, as a consequence, increasing effort is a sign of failure rather than success. Most of the data we have on phishing, such as reported in [22, 23, 17] measure activity rather than dollars. The popular (and indeed academic) press often presents measured increases in phishing activity as evidence of the success and profitability of the endeavor. As fishing grounds deplete

trawlers must go farther afield and stay at sea longer, but this is a sign of the poverty of the industry not its health. Similarly, we contend that increased activity results from the large number of phishers who try harder and harder to find the enormous returns that they believe to be there, but which earn them only their opportunity cost income.

Fourth, the easier phishing gets the lower the total revenue $R_{tot}(E)$. For example, suppose that a more automated way of phishing reduces the cost per unit of effort from a to b . The effect of this is shown in Figure 1. At the moment the innovation is introduced we have equilibrium at $R_{avg}(E) = a$. The innovators enjoy an average revenue of a for cost of b and make a profit while this persists. This is a short-lived opportunity however: those with cost b increase their effort to maximize the profit. The overall effort increases and a new equilibrium is established at $R_{avg}(E) = b$. This means that phishers from locations with a lower cost structure (*e.g.* alternative opportunities do not pay well) drive out those from more expensive locations. Thus the total return is determined by the alternative opportunities available to the least skilled people able to accomplish the task. If those alternatives are not good (*i.e.* phishers can find only meagre compensation elsewhere) then the total dollars lost due to phishing will be low. Further, since technology over time decreases the skill level needed for a task, the cost per unit of effort a decreases. This causes the equilibrium to be reached at higher levels of effort, and hence lower total revenue. Thus the total dollars lost due to phishing are probably decreasing over time. Increasing the sustainable harvest would require increasing the cost basis of the participants or restricting access to the resource pool.

Fifth, the economic equilibrium should be reached when $R_{tot}(E) = C_{tot}(E)$. This assumes that each decision maker behaves as a rational actor: when his revenue $R_{avg}(E)$ equals his cost he is indifferent to entry and exit from the phishing business. Thus, if he makes less than his opportunities elsewhere, he stops phishing. This is a somewhat simplified model. Economists have searched for examples where the shared resource yields either more revenue (*i.e.* $R_{avg}(E) > a$) or less (*i.e.* $R_{avg}(E) < a$) than the theory predicts. An example, for a time, of more revenue than predicted appears to be the example of English privateers who pillaged Spanish and French ships in the period 1625-1630 [7]. The conclusion reached is the considerable danger involved, and imperfect information (which consistently underestimated the quantity of bullion that Spain was bringing from her colonies) discouraged open access to the market. The excess profits were driven out within four years however. On the other hand there are at least three cases when average revenue is less than cost; *i.e.* people persist in trying to harvest the open access

resource even when they give up better opportunities elsewhere. This occurs when any of the following factors are involved:

- Emotional ties
- Gambling
- Poor information

Emotional ties obviously play a large role in an industry like fishing: a third generation fisherman who has spent his life at sea is less likely to leave even if he can earn more elsewhere. People consistently make economically irrational choices when gambling is involved: they persist in an activity where their costs exceed their revenue if they think that they are only one lucky break away from getting rich. For example, in an unregulated gold rush the average return and the best return are very different. Each prospector dreams of the latter, but most get the former. Finally, people can persist in uneconomic activities if they are misinformed about their true prospects. For example, while their current costs may exceed revenue they might believe that this will improve with time. We suggest that emotional ties play no role in phishing; *i.e.* few phishers are so tied to the skill and craft of their calling that they persist if it is uneconomic. However, it is likely that gambling and poor information both play some role. When the job appears to involve merely harvesting free money it seems safe to believe that many phishers dream of “hitting the jackpot.” Equally, many phishers might believe that their revenue will improve. Suppose, as we suggest, that the tragedy of the commons is at work, and that the total losses due to phishing are lower than generally reported and are dropping rather than rising with time. Each phisher knows his revenue and his cost. Suppose, for some phisher, these are $R_{avg}(E) - \epsilon$ and a , so he’s making a small loss in an opportunity cost sense. The rational economic choice is to stop phishing and do something else. But consistent reports of “easy money” may encourage him to think that he’s doing something wrong and that his returns will improve with time.

Sixth, the lack of information about the true returns effects the equilibrium. Phishers may persist longer than they should even if they make a loss. More pronounced, however, is the fact that consistent overstatements of revenue guarantees a steady stream of hopeful but mis-informed new entrants. In one sense this is good, and in another bad. It is good in that the new entrants increase the total effort and thus reduce $R_{tot}(E)$. This is especially so if (as newbies) they are willing to earn less than their costs while they learn. In doing so they push the total revenue further and further down, so the dollars lost decrease. This would mean that equilibrium would be achieved at $R_{tot}(E) < C_{tot}(E)$ and we would have the rather pleasing irony that phishers as a

whole would be losing money. While amusing, this situation is bad, in the sense that reducing the money stolen through phishing is not our only goal, which brings us to our final point.

Phishing is not solely (or perhaps even mainly) a problem of stolen dollars. If it were, and we wish to reduce $R_{tot}(E)$ we might simply encourage as many people as possible to phish. As each of them seeks to maximize their revenue they will drive the total return from the commons down. This isn't an interesting solution however. In reality the erosion of trust in email and web commerce is more significant than the lost dollars.

4. OBJECTIONS AND POSSIBLE PROBLEMS

A number of objections can be raised to this analysis. Chief among them are objections to the economic model and claims that reports show that there is a lot of money in phishing. We examine these in turn.

4.1 Objection: “The economic Model is too Simple, Phishing isn't like Fishing”

Numerous objections can be made to the economic model we have used. It can be argued that

- This is an equilibrium analysis, phishing is too new to be in equilibrium
- Phishing is an illegal activity and this affects entry and exits from the field
- Phishers' costs are not linearly related to effort
- Some phishers have very sophisticated operations that require little effort.

Each of these objections has merit, but none affects the analysis. The overall point is that when an activity offers returns superior to the alternatives it continues to attract newcomers until the returns are no longer superior. Unless phishing is governed by a set of Economic laws different from other human endeavors the invisible hand of the market drives the average earning to the opportunity cost.

The case of piracy [7] indicates that equilibrium can be reached in only a few years, even when information flow is poor, and greater danger and risk are involved. Indeed history teaches that lucrative opportunities do not lie unexploited long. News of a gold strike in the Klondike reached Seattle on July 17, 1897 and more than 100000 attempted the difficult trek to Dawson City in the following six months. By the Summer of 1899 the gold was substantially exhausted and the rush over.

The illegality of the activity certainly makes some reluctant to become involved, but this has no influence on earnings so long as there is a sufficient supply of recruits willing to exploit the opportunity. The case of low-level drug dealers confirms that unskilled work

seldom pays well even when it is illegal and dangerous [26].

It is certainly the case that phishers costs are not linearly related to efforts, and there can be a large difference between the efficiency with which different phishers pursue their victims. The straight line relation between costs and efforts $C_{tot}(E) = a \cdot E$ is a simplification, but a more complex relation doesn't change the outcome. The two key assumptions of the model are that revenue $R_{tot}(E)$ eventually falls with increasing effort, while costs $C_{tot}(E)$ are monotonically increasing with effort. It does not matter whether effort E is measured in hours or any other units: the effort expended keeps increasing so long as it is profitable, and the the revenue keeps falling. Equilibrium is reached when incentives to leave and incentives to enter the pool are in balance (cost and revenue are equal).

The argument that phishers have efficient automatic operations that require little effort merely argues that costs are very low.

4.1.1 *Maybe we're just on the early part of the curve?*

A natural objection is to question if we are really on the right part of the sustainable harvest curve in Figure 1. The tragedy of the commons (whereby increasing effort results in decreasing harvest) happens only when $E > E_x$. Is it not possible that we have not reached optimal yield yet (*i.e.* $E < E_x$)? It is impossible to reach equilibrium at $E < E_x$ unless the whole endeavor is impossible. This is so, since increasing effort gives increasing return; so new arrivals will continue to find the opportunity profitable. Thus equilibrium is not achievable at any $E < E_x$ unless $R_{tot}(E_x) < C_{tot}(E_x)$ or there is a barrier to new entrants. Since Gartner [13] estimates that 66% of the population had received phishing emails it is hard to argue that there's a large uncontacted population that represents a profitable opportunity.

4.1.2 *What about the Sinusoidal Predator-Prey Population Dynamics Model?*

The population dynamics of interacting predator prey species is sometimes modeled using the Lotka-Volterra equations. The solution gives that both populations oscillate sinusoidally, but with the predator population lagging the prey by 90° . Why does this model not apply? The reason essentially is that the Lotka-Volterra solution assumes a closed ecosystem and that the predator population grows and shrinks with births and deaths only. A large increase in available prey results in an increase in predators, but only slowly. By contrast the tragedy of the commons model assumes that new predators enter the system when the opportunities are good and leave when they are bad.

4.2 Objection: “What About all the Data Showing that Phishing Losses are Huge?”

The short answer is that data showing that phishing losses are huge crumble upon inspection. We review the main surveys and technical studies of phishing rate and losses in Section 5. We examine different estimates of the phishing rate (*i.e.* percent of the population who are phished each year) and the loss (*i.e.* the amount lost per victim).

4.2.1 Estimating Rate: Surveys versus Measurements

Most of the data are from victim surveys [12, 13, 14, 8, 9, 18, 29]. While surveys are a very valuable source of data, crime researchers have known for some time that victim surveys have several sources of bias:

- Selection bias (*i.e.* failure to contact a representative sample of the overall population)
- Refusal rate (*i.e.* rate at which contacted population refuse to respond to the survey)
- Telescoping (*i.e.* tendency of respondents to “throw in” incidents that do not fall within the time frame of the survey)
- Forgetting (*i.e.* tendency to omit crimes that do fall in the time frame or have been forgotten)
- Exaggeration of losses (*i.e.* tendency of victims to overstate rather than understate the magnitude of the wrong they have suffered).

Each of these can have significant effects on the outcome of a survey. Selection bias is potentially a very large factor. First, there is no registry of online users, so contacting a random subset of online users is exceedingly difficult. Phone surveys generally randomly select from registries of landline numbers and thus miss the cellphone-only population. Postal mail surveys tend to miss those who move often. Email is worst of all for performing a phishing survey since it would appear necessary to use the same technique that phishers use: send bulk mail to lots of addresses and hope for responses. The FTC [8, 9] and Javelin [18] surveys were done by phone. Gartner [12, 13, 14] does not specify their contact methodology.

Survey scientists have long known that achieving a low refusal rate among those contacted is vitally important to ensure that randomness in the contacted sample is carried over into the achieved sample. This is vitally important, since a high refusal rate amplifies any difference in response rates between victims and non-victims. For example, a contacted representative population contains victims and non-victims: $C = V + N$. If everyone responded we would estimate the victimization rate as $V/(V + N)$. But if only a fraction V_r and N_r of victims

and non-victims respectively respond we estimate the rate as

$$\frac{V \cdot V_r}{V \cdot V_r + N \cdot N_r} = \frac{V \cdot (V_r/N_r)}{V \cdot (V_r/N_r) + N}$$

When the overall victimization rate is low (*i.e.* $V \cdot (V_r/N_r) + N \approx V + N \approx N$) any difference in the victim and non-victim response rates enormously influences the estimate. This can be very pronounced when the response rate is low; *e.g.* if $V_r = 5N_r$ (victims are $5\times$ more likely to respond) then the estimated victimization rate is almost $5\times$ the true rate. This effect can be so large that it is regarded as good practice in victim surveys to follow up with non-respondents (*i.e.* those who refuse to participate) and ask whether they were victims or not, even if they do not answer more detailed questions. This makes it possible to estimate whether victims and non-victims are responding at different rates, and if so, adjust for the bias. The response rate for the FTC phone survey [9] was 26%, and the response rates on email surveys can be an order of magnitude lower. Thus all of the ingredients for a very biased estimate are present in each of the surveys: difficulty contacting a random sample, high refusal rate, low victimization rate and greater likelihood that victims respond.

Rather obviously, the phishing victimization rate is small. In fact, in all of the victim surveys except [14] the margin of errors for the 95% confidence interval is larger than the estimated phishing rate. We tabulate the margins of error in Table 1. Observe that even though the Javelin 2005 [18] and Gartner 2005 [12] produce phishing estimates for the same period that differ by almost an order of magnitude (*i.e.* 0.07% and 0.5%) that is still well within the margin of error. It is very misleading to state (as Gartner does [14]) that “phishing attacks in the United States soared in 2007” on the basis of an increase (from 2006) that is less than the margin of error.

Being free of these biases the rate measurements performed by Florêncio and Herley [10] and Moore and Clayton [28] are likely far more accurate than any of the surveys. The fact that their estimates of the phishing victimization rate, using entirely different measurements, agree so well (0.4% and 0.34%) encourages us to suggest that the true rate is somewhere in this neighborhood. The surveys biases mentioned can comfortably account for difference with the rate estimated by Gartner.

4.2.2 Estimating Dollar Losses and Mean vs. Median

In all of the surveys examined in Section 5 victims are self-reporting losses. This is problematic. Indeed a few sanity checks reveal that the self-reported \$47bn in ID theft losses [8] is almost certainly enormously exaggerated. By way of benchmarking, the total reported 2003 profits of the top five banks in the US (Citi, BoA,

HSBC, WellsFargo and JP Morgan Chase) was \$59bn. Alternatively, \$50 billion would give an income \$500k each annually to a population of 100k identity theft professionals (which is approximately four times the number of cardiologists in the US).

In some of the surveys (*e.g.* [8, 18, 9]) the victim was also asked how much they think the thief obtained. In some (*e.g.* [12, 13, 14]) the victim was asked how much of their loss they eventually recovered. For some categories of identity theft in [8] there is almost an order of magnitude between the amount that victims were out of pocket and the amount that they believed the thief obtained (see responses to Q29 and Q30 in [8]). This leaves open the possibility that victims exaggerate the amount lost, and/or underestimate the ability of banks to halt and reverse transactions or recoup funds. Most of the phishing surveys do not make clear whether the victim was asked for out-of-pocket loss, or what they thought the thief gained. Thus, from [8], we have an order of magnitude uncertainty about these numbers.

Almost all of the surveys take the average reported loss per victim and scale this loss to the entire population. It cannot be emphasized strongly enough that a few victims who claim extravagant losses can bias the *average* numbers greatly upward (the same is not true of the median). This error is not symmetric since victims who understate losses cannot exert similar influence. When exaggeration occurs the enormous potential to influence the results is made clear in [9]. Footnote 8 of [9] states that two individuals reported losses of \$999999 and \$485000. For the first examination of the interview indicated that the crime claimed was not ID theft. For the second “the record seemed inconsistent with the loss of this much money.” Both losses were excluded from the calculation. Had the \$485k loss been included, the estimated total loss would have increased from \$15.6 billion to \$27.9 billion. *Thus, a single individual who exaggerated made almost a factor of two difference in the estimated loss.* The authors of [9] point out that the apparent drop in ID theft losses from \$47 billion in [8] to \$15.6 billion in [9] may be due this methodology change (*i.e.* excluding data from individuals who claimed extravagant losses that did not withstand scrutiny). Clearly, without sanity checking a small number of individuals can exert enormous bias on the average. There is, of course, a possibility that a small number of individuals suffer great losses, but just as [9] found, some also exaggerate and misreport. In either case, if a very small number of respondents account for the bulk of the losses the dollar estimates must be regarded as very noisy. Probably for this reason [9] cites median losses and no longer gives averages.

Following [9] Gartner in 2007 [14] shows a factor of almost 4.5 between the median and mean loss (\$200 and \$886 respectively). This is an extraordinary gap, and

strongly suggests that a small number of respondents are influencing the mean (just as in [9]).

4.3 Objection: “Phishers would not be doing this if it weren’t profitable”

If the Economics of phishing are as dismal as they appear then why do phishers persist? Why would anyone keep phishing if they were losing money? Observe that phishing can remain a going concern even if no individual persists so long as new entrants keep arriving to replace those leaving the business. Recall, would-be new entrants to the phishing business have access only to the “easy money” public stories everyone else sees. Stories of “instant riches” and “easy money” dependably attract newcomers whether to a gold rush or to phishing.

Indeed one explanation of the thriving trade in phishing related services reported in [23, 17] is that phishers with more experience prey upon those with less. That is, those who have tried phishing and found it unprofitable or marginally profitable find it better to sell services to those who haven’t reached that conclusion yet. A resource (*e.g.* the use of a botnet to send spam or servers to host the site) can be rented out for more than it will yield if the buyer overestimates the likely return. And a constant supply of “easy money” stories ensures that new entrants overestimate the expected returns. Ford and Gordon [11] suggest that attacks on the business models of malware can succeed where technology alone cannot. Thus ensuring better information might decrease the flow of new entrants to the phishing business.

4.4 Objection: “OK, so what’s your estimate?”

Having questioned the accuracy of a number of other estimates it seems only fair that we advance one. On rate we use [10] and [28] to estimate that 0.37% of web-users are phished annually. By this we mean that they type their passwords at phishing sites. It does not follow that all lose money, or even have their accounts compromised. Some phishing servers will be seized before the credentials are harvested, some users will realize their mistake and change password, some accounts may have no money or not be enabled for online transfers, and some banks may spot the attempt at fraud before any transfer happens. To account for all of these factors we estimate that half of those who type their password at a phishing site have their account compromised. If we assume that all of these victims lose the median found in [14] we get annual US losses of $0.0037 \times 0.5 \times 165e6 \times 200$ or \$61 million annually. It may be puzzling to use the median rather than the mean of [14], but the 4.5× factor between them indicates that the reported mean cannot be trusted (see Section 4.2.2).

We emphasize that this estimate might easily be off by

a factor of two or more. While, by no means small, this is a factor of 50 lower than [14]. It is of course hard to know how many phishers share this harvest. Franklin *et al.* observed 113k unique participant IDs in their study. A \$61 million harvest divided 113000 ways would indeed suggest that the average phisher earns hundreds rather than thousands of dollars for his efforts.

5. RELATED WORK

5.1 Phishing Surveys

There are several surveys estimating phishing activity, but it is difficult to compare the results of these estimates. With the exception of [8] only a summary is available in most cases. Surveys often group losses with other types of crimes, unrelated to phishing (e.g., identity theft by a family member). When a dollar amount for loss per victim is given it is often unclear whether this is what the respondent claims to be their out-of-pocket loss, or what they believe the thief received (where both were asked in [8] there is a large difference). In some cases, *e.g.* [12, 13], one figure is quoted for amount lost, and another for amount recovered. It is unclear whether the amount recovered is due to the bank retrieving the victim's funds from the thief or whether the bank is absorbing that portion of the loss (it makes no difference to the victim, but does affect the amount the phisher receives). Having said that, let us list the findings of the principal available surveys.

5.1.1 FTC Survey 2003 [8] and 2007 [9]

A widely cited survey sponsored by the Federal Trade Commission in 2003 did a phone survey of 4057 US adults and asked a variety of questions related to identity theft and fraud. This study is the source of the oft-quoted "\$47 billion in identity theft losses." While this survey did not address phishing directly it found the rate all types of fraud on existing accounts to be 0.7% with an average loss of \$2100.

A similar survey in 2006 [9] of 4917 US adults inferred losses of \$15.6 billion. As noted in [9], and discussed in Section 4.2.2 the apparent drop with respect to [8] may be entirely due to methodology changes.

5.1.2 Truste 2004 survey [29]

Truste found a 2% phishing victim rate among 1335 surveyed adults. Their \$500 million national loss implies an average of \$116 per victim.

5.1.3 Javelin 2005 survey [18]

Javelin in 2005 surveyed 4000 consumers by phone and used methodology similar to [8]. A total of 4.25% said they had been victims of ID theft (all types) in the last year. They go on to say 1.7% of the victims were phishing victims with an average loss of \$2820/victim.

But 1.7% of 4.25% of 4000 people is only 3 people, so the dollar estimates are extremely noisy. They extrapolate to "\$367 million phishing losses/yr in the US."

5.1.4 Gartner 2005 [12], 2006 [13] and 2007 [14]

Gartner did a survey published in June 2005 of 5000 web consumers. They found 0.5% of respondents claimed to be victims of phishing. They extrapolate to get "\$929 million phishing losses/yr in US, among 1.2 million victims." This is an implied loss of \$774/victim. They repeated the study in 2006 [13] and found a loss of \$1244 per victim among 2.25 million victims (or 1.05% of a 215 million adult US population). In 2007 [14] they found 2.18% of respondents claimed to be victims with an average loss of \$886, but a median loss of \$200. In common with [9] there is a large difference between the median and mean loss per victim. We discuss this difference in Section 4.2.2. Victims reported recovering 80%, 54% and 64% of their losses in 2005, 2006 and 2007 respectively.

5.2 Technology Estimates

5.2.1 Password Re-Use Study [10]

Florêncio and Herley report the findings of measuring Password Re-Use among users of the Windows Live Toolbar. Among 436k users, over a three week period, they found that 101 entered previously-used passwords at sites on a phishing blacklist, leading to an annualized phish victim rate of 0.4%. There is a bias toward more active users (since they downloaded general purpose toolbar), but the study involves many more users than any of the surveys above, and measures what users actually do, rather than what they remember and say they did.

5.2.2 Phishing Site Takedown Study [28]

Moore and Clayton did a survey of phishing sites from a live feed. They estimated the number of victims per site using the ingenious observation that the popular reporting package Webalizer is running on many phishing sites. They estimate (excluding rockphish gang sites) 9347 phishing sites over a year, and 30 victims per site, giving 280k victims annually. They double this to account for the approximately 50 % market share enjoyed by the rockphish gang. This would give 0.034% of users phished annually based on a population of 165 million online users in the US.

5.2.3 Simulated Phishing Attack Study [19]

Jakobsson and Ratkiewicz simulated a real phishing attack by emailing 237 users with a phishing link and then measuring the number of responses. Four different experimental scenarios were tried. Their results suggest that a given phishing attack can have a yield of about 11% of users. However this measurement is only of users

clicking on a link and not of entering any information or actually having money stolen.

5.2.4 *Underground Economy Study* [17]

Franklin *et al.* measured traffic on one of the underground economy information markets. They detail an enormous volume of traffic advertising to buy and sell stolen login credentials and credit card numbers. However, they have no means of observing actual transactions and hence form no estimate of phishing losses. Their tally of the account balances of purportedly compromised accounts over the seven month period was \$54 million. These are the face-value claims of the would-be sellers of the credentials which are likely to be overstated.

6. CONCLUSION

We have advanced a view of the economics of phishing that challenges accepted wisdom. Far from being an easy money proposition we claim that phishing is a low skill, low reward business, where the average phisher makes about as much as if he did something legal with his time. The absence of data documenting large phishing gains suggests that this view has merit. We find that the data from widely cited victim surveys [12, 13, 14, 8, 9] are noisier and more biased than is generally realized.

It is interesting to wonder why the Gartner and FTC estimates are repeated without scrutiny when they appear noisy at best. Shafer [5] suggests that the answer lies in two classic papers by Singer [27] and Reuter [24]. Singer examined the commonly accepted numbers for heroin-related crime in NYC and found that, while they failed to survive basic sanity checks, they were nonetheless much quoted. Another amusing example of the trajectory of a quotable but poorly sourced claim (that internet traffic was doubling every hundred days) is examined by Odlyzko [6]. Reuter revisited Singer's problem and postulated that mythical numbers circulate without criticism when there is a constituency with an interest in having the reported numbers be high, but no constituency with an interest in having those numbers be accurate, and an absence of scrutiny from academic researchers. Phishing estimates would appear match those criteria.

We think that this economic analysis has important implications in addressing the problem on a macro level. If we are correct that large phishing dollar losses are an exaggeration, an important conclusion is that repeating those claims "feeds the beast," perpetuates the myth of the infinitely capable superuser attacker [21], and attracts poorly-informed new entrants to phishing. While it drives the dollar losses down further, it pollutes the ecosystem with yet more spam.

Finally, we would like to emphasize and re-emphasize that, even if the dollar losses are smaller than often

believed, we believe that phishing is a major problem. There are many types of crime where the dollars gained by the criminal are small relative to the damage they inflict. This appears to be the case with phishing. If the dollar losses were zero the erosion of trust among web users, and destruction of email as a means of communicating would still be a major problem.

Acknowledgements: the authors would like to thank Jan Feyereisl, Bob Blakley, Joshua Ladau and Christian Probst for detailed comments that helped improve the presentation.

7. REFERENCES

- [1] <http://ha.ckers.org/blog/20070508/phishing-social-networking-sites/>.
- [2] <http://en.wikipedia.org/wiki/Overgrazing>.
- [3] http://www.darkreading.com/document.asp?doc_id=116574&f_src=darkreading_section.296.
- [4] <http://fisherieseconomics.googlepages.com/openaccess>.
- [5] <http://www.slate.com/id/2144508/>.
- [6] A. Odlyzko. Internet traffic growth: Sources and implications. *Proceedings of SPIE*, 2003.
- [7] J. Conybeare and T. Sandler. State-sponsored Violence as a Tragedy of the Commons: England's Privateering Wars with France and Spain, 1625-1630. *Public Choice*, 1993.
- [8] Federal Trade Commission. Identity Theft Survey Report. 2003. <http://www.ftc.gov/os/2003/09/synovatereport.pdf>.
- [9] Federal Trade Commission. Identity Theft Survey Report. 2007. www.ftc.gov/os/2007/11/SynovateFinalReportIDTheft2006.pdf.
- [10] D. Florêncio and C. Herley. A Large-Scale Study of Web Password Habits. *WWW 2007, Banff*.
- [11] Ford R., and Gordon S. Cent, Five Cent, Ten Cent, Dollar: Hitting Spyware where it Really Hurt\$. *NSPW*, 2006.
- [12] Gartner. Identity Theft Survey Report. 2005. http://www.gartner.com/press_releases/asset_129754_11.html.
- [13] Gartner. Phishing Survey. 2006. <http://www.gartner.com/it/page.jsp?id=498245>.
- [14] Gartner. Phishing Survey. 2007. <http://www.gartner.com/it/page.jsp?id=565125>.
- [15] H. S. Gordon. The Economic Theory of a Common-Property resource: The Fishery. *Journal of Political Economy*, 1954.
- [16] G. Hardin. The Tragedy of the Commons. *Science*, 1968.
- [17] J. Franklin and V. Paxson and A. Perrig and S. Savage. An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants. *Proc. CCS*,

Source	Webusers phished annually	Margin of error	Average loss per victim	Median loss per victim
FTC 2003	< 0.7%	1.5%	\$2100	NA
FTC 2007	< 1.5%	1.4%	NA	\$0
Truste 2004	2.0%	2.7%	\$125	NA
Javelin 2005	0.07%	1.5%	\$2820	NA
Gartner 2005	0.5%	1.4%	\$774	NA
Gartner 2006	1.05%	1.4%	\$1244	NA
Gartner 2007	2.18%	1.4%	\$886	\$200
Florêncio and Herley	0.4%	0.6%	NA	NA
Moore and Clayton	0.34%	NA	NA	NA

Table 1: Estimates of phishing rates and losses per victim from various surveys. Note that FTC 2003, FTC 2007 and Javelin 2005 as phone surveys are estimating rate over the adult US population (215 million [8]), while all the others estimate over the online adult population (estimated variously at about 165 million). Also FTC 2003 and FTC 2007 did not ask about phishing directly.

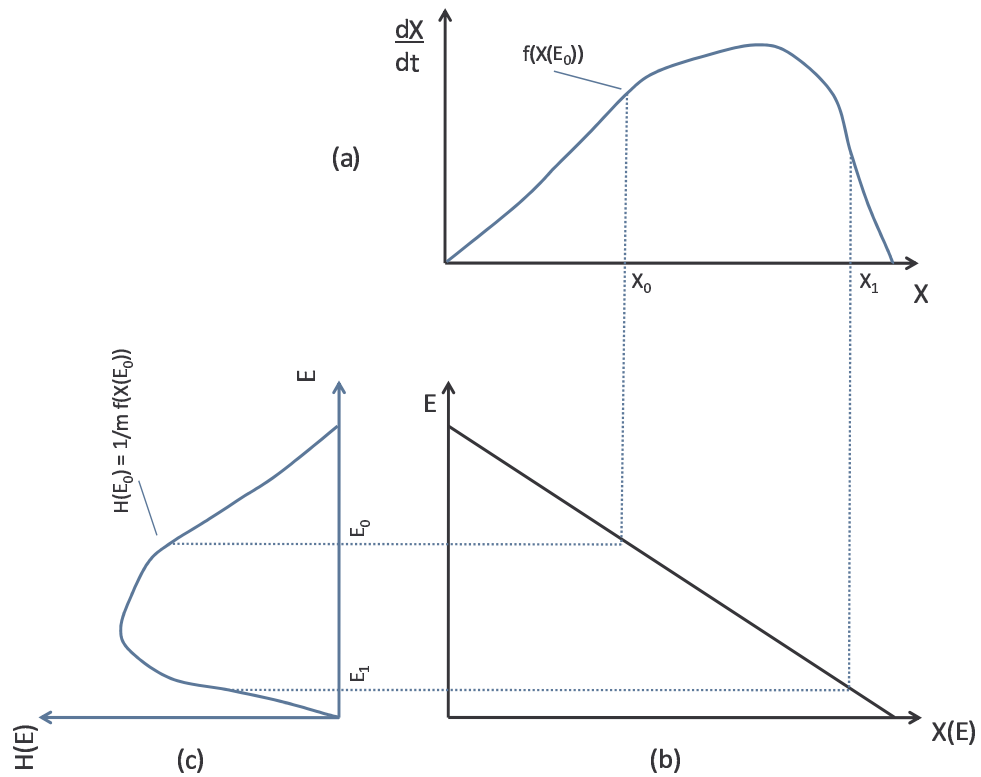


Figure 2: Construction of the sustainable harvest curve. (a) The number of phishable dollars added per unit time grows with the size of the pool, but then growth slows and drops to zero (the pool is finite) (b) The pool of phishable dollars depends on the effort E : for no phishing effort it achieves it's maximum, at some level of effort there are no phishable dollars left. (c) The sustainable harvest curve. For a fixed effort, e.g. E_0 , we determine the phishable dollars $X(E_0)$ from (b) and hence $f(X(E_0))$ from (a). Since, in equilibrium $H(E_0) = 1/m \cdot f(X(E_0))$ we get the sustainable harvest curve by repeating for many values of E .

2007.

- [18] Javelin. Identity Theft Survey Report. 2003.
http://www.javelinstrategy.com/uploads/505.RF_Phishing.pdf.
- [19] M. Jakobsson and J. Ratkiewicz. Designing Ethical Phishing Experiments: A Study of (ROT13) rOnl Query Features. *Proc. WWW*, 2006.
- [20] P. A. Neher. The Pure Theory of Muggery. *American Economic Review*, 1978.
- [21] P. Ohm. The Myth of the Superuser: Fear, Risk, and Harm Online. *UC Davis Law Review*, 2008.
- [22] Anti-Phishing Working Group.
<http://www.antiphishing.org>.
- [23] R. Thomas and J. Martin. The Underground Economy: Priceless. *Usenix ;login;*, 2006.
- [24] P. Reuter. The (continued) Vitality of Mythical Numbers. *Public Interest*, 1987.
- [25] S. Iudicello and M. Weber and R. Wieland. Fish, Markets and Fishermen: the Economics of Overfishing. *Island Press*, 1999.
- [26] S.D. Levitt and S. J. Dubner. Freakonomics: A Rogue Economist Explores the Hidden Side of Everything. *William Morrow*, 2005.
- [27] M. Singer. The Vitality of Mythical Numbers. *Public Interest*, 1971.
- [28] T. Moore and R. Clayton. Examining the Impact of Website Take-down on Phishing. *Proc. APWG eCrime Summit*, 2007.
- [29] TrustE. Phishing Survey. 2004.
http://www.truste.org/about/press_release/09_29_04.php.