

How Much Security is Enough to Stop a Thief?

The Economics of Outsider Theft via Computer Systems and Networks

Stuart E. Schechter and Michael D. Smith

Harvard University
{stuart,smith}@eecs.harvard.edu

Abstract. We address the question of how much security is required to protect a packaged system, installed in a large number of organizations, from thieves who would exploit a single vulnerability to attack multiple installations. While our work is motivated by the need to help organizations make decisions about how to defend themselves, we also show how they can better protect themselves by helping to protect each other.

Keywords: Security Economics, Threat Models, Theft, Exploits

1 Introduction

Before deploying a new laptop computer or installing a new email program, a prudent organization will want to ensure that that system provides adequate security. An organization can determine the security of a computing system by measuring the cost of finding and exploiting a security vulnerability in that system [1]. This measure, known as the *cost to break*, is most effective when you also know how much security your organization requires. To answer the question of how much security is enough, you must first determine what types of adversaries you may need to defend against and what choices are available to each type of adversary.

To this end we introduce *economic threat modeling* as a tool for understanding adversaries motivated by financial gain. Specifically, we model those thieves outside the target organization who would enter via an unreported vulnerability in one of the target's *packaged systems*, those systems that are replicated and installed in many organizations. This model can then be used to estimate what these thieves are willing to pay for system vulnerabilities and how secure the system needs to be to make theft an unprofitable proposition.

Standardization and the wide-spread use of packaged systems means that an organization must look outside itself to determine how thieves might exploit the packaged systems it uses. An organization cannot consider itself safe simply because it would cost a thief more to find and exploit a new vulnerability in one of the organization's systems than that thief could gain by exploiting the vulnerability. Instead, the organization must consider that the thief evaluates his potential financial gains in a global context; he decides whether it is profitable to find and exploit a new vulnerability based on all of the organizations deploying

that packaged system. Unless an organization can afford to build every system it uses from scratch, it must also measure its security in a global context.

As this paper demonstrates, the global nature of security also works to the advantage of the organization, since the community of all potential victim organizations has similar interests in thwarting attacks. For example, we argue that one way organizations can lower the damage resulting from unreported vulnerabilities is by sharing information about recent attacks. Monitoring companies and federally funded centers such as CERT can help organizations to collect and disseminate such information.

We make a case for the construction of economic threat models in Section 2. We define outside theft in Section 3 and introduce two subcategories. Serial theft, in which a criminal makes repeated uses of the same exploit to rob one victim after another, is analyzed in Section 4. Parallel theft, in which thieves automate the process of exploiting vulnerabilities in order to attack many systems at once, is addressed in Section 5. Applicability and relevance of our work to other threats is discussed in Section 6. In Section 7 we look at the implications of the results to all the members of the defense. Related work follows in Section 8, and we conclude in Section 9.

2 Economic threat models

An economist observing a fisherman and wondering why it is that he fishes might pose the following two questions: How difficult it is for the man to catch fish, and how much are consumers willing to pay for fish? Economic threat models are designed to answer these same basic questions. The fundamental importance of these two questions to understanding security becomes clear when one looks at the fisherman and the consumer from a new perspective—that of a fish.

As the security of the fish depends on the number of people who choose to fish and the resources (rods, lines, nets) at their disposal, the security of a system depends on the number of people who stand to profit from attacking it. As in fishing, the choice to attack depends on what one stands to gain given the costs and resources available.

Traditional threat models help us understand who the adversary may be and what motivates them, but do so in a qualitative, not quantitative manner. To understand where these models fall short, it is important to understand where they fit into the security process. Figure 1 shows a representation of this process, traditionally separated into the steps of prevention, detection, and response, expanded to better detail the prevention process. Working backwards, we see that in order to make the right-sized investment in security we must be able to determine the desired level of security. By quantitatively determining the point at which the costs to a potential attacker outweigh the benefits of attack, we can identify this desired security level. Traditional threat models fall short because they do not provide a quantitative measure of much security is enough to deter a given adversary.




Fig. 1. An illustration of the security process with an emphasis on protective steps.

3 The threat of outside theft

One cannot determine how much security is required to protect against adversaries without making assumptions about who the adversaries are, what motivates them, and what resources are at their disposal. For the next several sections, we focus on thieves outside an organization. *Outside thieves* are individuals or groups of individuals not affiliated with your organization who attack your systems for financial gain. We assume that thieves, in contrast to adversaries such as terrorists, behave rationally in so far that they will not stage attacks that are expected to lead to their own financial loss.

Thieves have many ways to profit from exploiting system vulnerabilities. The most obvious method is for the thief to steal information, such as trade secrets, customer lists, credit card numbers, or cryptographic keys, and then sell that information to others. However, a thief may not actually need to take anything during an attack to create a situation where she can sell something for profit. For example, the thief may create back doors in the systems he attacks and then later sell that access to the highest bidder. Alternatively, the thief may only change the state (e.g. a bank balance) on the target machines. Young and Yung [2] describe a state-changing attack that involves encryption of the data on the victim's machine; the thief profits by ransoming the decryption key.

3.1 Serial thieves

Serial thieves exploit an unreported vulnerability in a packaged system to attack victim after victim. By concentrating on one (or a small number) of victims at a time, the serial thief can carefully survey the value and location of what each victim is protecting and then maximize the loot obtained while minimizing risk.

A serial thief's crime spree ends when he is caught, when the vulnerability he has learned to exploit has been detected and patched on all target installations, or when the reward of committing another theft with this exploit no longer outweighs the risk of losing the loot already collected.

3.2 Parallel thieves

Parallel thieves automate their attack to penetrate a large number of targets at the same time. Automation has the clear appeal of multiplicatively increasing the potential loot beyond that obtainable by a more meticulous serial thief in

the same time period. In fact, this may be the only practical way to rob a large number of victims if exploiting the vulnerability is likely to reveal it to the defense and the manufacturer of the vulnerable system is quick to release patches.

There are, however, three costs to automation that relate to our later analysis. First, automation requires the thief to be able to identify the loot without direct human guidance. There is no doubt this can be done (e.g., by looking for files of a particular type), but a parallel thief will not be able to customize the attack to ensure that the maximum value is captured from each victim. Second, automation also forces the thief to make a fixed set of common assumptions about the vulnerability and its surrounding defenses. If the assumptions are incorrect, the attack may fail for a particular organization or worse, may significantly increase the chance that the thief is discovered and later convicted. Finally, it is harder to hide a flood of automated penetrations occurring in a relatively short time period, and many intrusion detection tools look for such statistically significant abnormalities in system usage. This is compounded by the fact that as the number of victims increases, so does the number of transactions required to collect payment for the stolen loot. As we discuss later, there is risk in each of these transactions.

4 Serial Theft

A target's attractiveness to a serial thief is the expected income to the thief from attacking it. This is in turn a function of the probability of success and the amount of loot that may be protected by the target. To simplify the initial analysis, we start by assuming that a thief will attack the most attractive targets, and that there is an unlimited number of homogeneous targets that are all equally attractive. We later extend the analysis by removing this assumption and modelling the unique qualities of each target.

4.1 Homogeneous Targets

We model the choices of the serial thief using a number of variables to represent properties of his environment. The motivation for each crime is the amount of loot that he expects to capture if the theft is successful.

ℓ	Loot, or the value obtained from a successful theft.
--------	--

For every theft, there is a chance that the thief will be caught, convicted, and punished. We thus define the probability of being convicted and attempt to quantify in dollar terms the value of the punishment.

\mathcal{P}_c	Probability of being caught, convicted, and punished for the theft.
-----------------	---

F	Fine paid by the thief if convicted, or the dollar equivalent of other punishment levied in the event of conviction, including the value of any confiscated loot.
-----	---

Catching and convicting a serial thief may be difficult, especially if the thief is in a foreign jurisdiction. Another way to stop a thief who is taking advantage of an unknown vulnerability in a system is to detect how he is breaking in. Each time the thief uses an exploit to break into a system, there is a chance that an intrusion detection system or monitoring firm will be able to observe the thief's means of attack. Once the vulnerability is discovered, intrusion prevention systems may be trained to detect attacks against it and the vulnerable system's manufacturer may distribute a patch. The serial thief will no longer be able to use this exploit to attack organizations that patch their systems.

\mathcal{P}_d Probability that use of the exploit will expose it to the defense and the vulnerability will be patched as a result.

We assume that if the thief is caught and convicted, his methods will be divulged. Thus $\mathcal{P}_c \leq \mathcal{P}_d$.

Finally, there is always a chance that a given attack will fail for reasons other than that the vulnerability was detected and a patch was put into place. Perhaps the target's intrusion detection system had learned just enough from reports of previous attacks at other targets to enable it to recognize a new attack. Perhaps a behavior-based intrusion detection system recognized the theft to be unusual activity and stalled while monitoring personnel were notified.

\mathcal{P}_f Probability that the attack fails, possibly because it is repelled before the criminal can reach the loot.

Note that \mathcal{P}_f and \mathcal{P}_c are independent. An attack may succeed but the thief may be caught and convicted; an attack may fail but the thief may not be captured or convicted; or an attack may succeed and the thief may escape capture or conviction.

To simplify our notation, the probability that the exploit will not be divulged is written $\overline{\mathcal{P}}_d = 1 - \mathcal{P}_d$. The probability that the attack does not fail is written $\overline{\mathcal{P}}_f = 1 - \mathcal{P}_f$.

The expected profit from the i th theft, assuming the exploit has not yet been detected and patched, is the expected loot, $\overline{\mathcal{P}}_f \ell$, minus the expected fine if convicted, $\mathcal{P}_c F$:

$$\overline{\mathcal{P}}_f \ell - \mathcal{P}_c F$$

The expected profit from the i th and all additional thefts is labeled $E_{i \rightarrow \infty}$. We account for the additional value of the future thefts by adding the expected value of those thefts on the condition that the exploit can be used again, $\overline{\mathcal{P}}_d E_{i+1 \rightarrow \infty}$.

$$E_{i \rightarrow \infty} = \overline{\mathcal{P}}_f \ell - \mathcal{P}_c F + \overline{\mathcal{P}}_d E_{i+1 \rightarrow \infty}$$

Expanding reveals a common recurrence of the form $x^0 + x^1 + x^2 + \dots$, which for $0 < x < 1$ is equal to $\frac{1}{1-x}$.

$$\begin{aligned}
E_{i \rightarrow \infty} &= (\bar{\mathcal{P}}_f \ell - \mathcal{P}_c F) \cdot (1 + \bar{\mathcal{P}}_d + \bar{\mathcal{P}}_d^2 + \bar{\mathcal{P}}_d^3 + \dots) \\
&= (\bar{\mathcal{P}}_f \ell - \mathcal{P}_c F) \cdot \frac{1}{1 - \bar{\mathcal{P}}_d}
\end{aligned}$$

Thus the value of an exploit to a serial thief attacking homogeneous targets is:

$$E = E_{1 \rightarrow \infty} = \frac{\bar{\mathcal{P}}_f \ell - \mathcal{P}_c F}{\mathcal{P}_d} \quad (1)$$

This result is quite revealing for the defense. It shows that even if you can't increase the chance of convicting a thief or even thwart attacks that haven't been seen before, you can cut the thief's expected revenue in half by doubling \mathcal{P}_d . That is, by doubling the probability that the vulnerability used in the attack will be revealed to the defense and subsequently repaired, the amount of loot that the thief expects to extract from society is reduced by half. Deploying security tools to thwart attacks is also an effective deterrent. Halving the probability that an attack succeeds at each target will also reduce the value of the exploit, by at least if not more than half.

When evaluating different approaches for changing the probabilities in Equation 1, an organization should remember that customized security tools (i.e., ones not readily available to the adversary) are the ones that will be most effective. Unfortunately, these tools are also the most expensive to implement. Security tools that are themselves packaged systems may not be as effective for this purpose. If such packages are available to the thief, he will have had every opportunity to customize the exploit tool in an attempt to circumvent the defense and to avoid detection during an attack. This is a strong argument for using a monitoring company that uses customized detection systems that are unavailable to the adversary.

4.2 Unique Targets

Instead of viewing all targets as homogeneous, we now assume that each is unique. Each potential target organization has two goals in formulating its defense: to minimize the profitability of attack, and to make itself the last target a thief would choose to attack. Action towards the first goal reduces the likelihood that thieves will attack. The second goal is motivated by the fact that if a string of thefts does take place, then the later the organization falls in the set of targets, the more likely it is that the spree will end before the organization is attacked.

We label each target t_i where t_1 is the first target attacked, t_2 the second, and so on. An ordered set of targets t_1, \dots, t_i is written \mathbf{T}_i . The loot for each target is defined as ℓ_i and the punishment for the i th attack is F_i .

$\mathbf{T}_i = t_1, \dots, t_i$	An ordered set of targets attacked.
----------------------------------	-------------------------------------

ℓ_i	The loot obtained by attacking target t_i .
----------	---

F_i	The punishment if caught and convicted after the i th attack.
-------	---

The probability that the vulnerability used to attack the system is detected is defined two different ways.

$\overline{\mathcal{P}}_d(t_i \mathbf{T}_{i-1})$	Probability that, if targets t_1 through t_{i-1} have been attacked, and the exploit has yet to be discovered, that it will not be discovered when t_i is attacked.
--	---

$\overline{\mathcal{P}}_d(\mathbf{T}_i)$	Probability that the thief can attack all targets, t_1 through t_i , in order, without the exploit being discovered and fixed.
--	--

The latter probability can be expressed inductively in terms of the former: the probability that all i attacks remain undetected is the probability that the first $i - 1$ attacks remained undetected multiplied by the probability that the i th attack remains undetected.

$$\begin{aligned}\overline{\mathcal{P}}_d(\mathbf{T}_i) &= \overline{\mathcal{P}}_d(t_i|\mathbf{T}_{i-1})\overline{\mathcal{P}}_d(\mathbf{T}_{i-1}) \\ \overline{\mathcal{P}}_d(\mathbf{T}_1) &= \overline{\mathcal{P}}_d(t_1|\emptyset) = \overline{\mathcal{P}}_d(t_1) \\ \overline{\mathcal{P}}_d(\mathbf{T}_0) &= \overline{\mathcal{P}}_d(\emptyset) = 1\end{aligned}$$

We also need functions to describe the probabilities that the thief will or won't be caught and that the attack won't fail.

$\mathcal{P}_c(t_i \mathbf{T}_{i-1})$	Probability that, if targets t_1 through t_{i-1} have already been attacked and the thief was not caught, the thief will be caught and convicted when he attacks t_i .
---------------------------------------	--

$\overline{\mathcal{P}}_f(t_i \mathbf{T}_{i-1})$	Probability that, if targets t_1 through t_{i-1} have already been attacked, the attack on t_i will <i>not</i> fail to produce loot.
--	--

The expected income from attacking unique targets is an extension of the same recurrence shown in the homogeneous target case.

$$\begin{aligned}
E_{1 \rightarrow n} &= \bar{\mathcal{P}}_f(t_1) \ell_1 - \mathcal{P}_c(t_1) F_1 + \bar{\mathcal{P}}_d(t_1) E_{2 \rightarrow n} \\
&= \bar{\mathcal{P}}_f(t_1) \ell_1 - \mathcal{P}_c(t_1) F_1 \\
&\quad + \bar{\mathcal{P}}_d(t_1) [\bar{\mathcal{P}}_f(t_2|\mathbf{T}_1) \ell_2 - \mathcal{P}_c(t_2|\mathbf{T}_1) F_2 + \bar{\mathcal{P}}_d(t_2|\mathbf{T}_1) E_{3 \rightarrow n}] \\
&= \bar{\mathcal{P}}_f(t_1) \ell_1 - \mathcal{P}_c(t_1) F_1 \\
&\quad + \bar{\mathcal{P}}_d(t_1) [\bar{\mathcal{P}}_f(t_2|\mathbf{T}_1) \ell_2 - \mathcal{P}_c(t_2|\mathbf{T}_1) F_2] \\
&\quad + \bar{\mathcal{P}}_d(t_1) \bar{\mathcal{P}}_d(t_2|\mathbf{T}_1) E_{3 \rightarrow n} \\
&= \bar{\mathcal{P}}_d(\mathbf{T}_0) [\bar{\mathcal{P}}_f(t_1|\mathbf{T}_0) \ell_1 - \mathcal{P}_c(t_1|\mathbf{T}_0) F_1] \\
&\quad + \bar{\mathcal{P}}_d(\mathbf{T}_1) [\bar{\mathcal{P}}_f(t_2|\mathbf{T}_1) \ell_2 - \mathcal{P}_c(t_2|\mathbf{T}_1) F_2] \\
&\quad + \bar{\mathcal{P}}_d(\mathbf{T}_2) E_{3 \rightarrow n}
\end{aligned}$$

The recurrence simplifies to the following summation where n is the number of thefts.

$$E = E_{1 \rightarrow n} = \sum_{i=1}^n \bar{\mathcal{P}}_d(\mathbf{T}_{i-1}) [\bar{\mathcal{P}}_f(t_i|\mathbf{T}_{i-1}) \ell_i - \mathcal{P}_c(t_i|\mathbf{T}_{i-1}) F_i] \quad (2)$$

As long as there are systems for which the thief's chosen vulnerability has yet to be patched, he will attack another target if the next term in Equation 2 is positive, increasing the expected profit. That is, attacking target t_{n+1} increases E so long as:

$$\bar{\mathcal{P}}_f(t_{n+1}|\mathbf{T}_n) \ell_{n+1} - \mathcal{P}_c(t_{n+1}|\mathbf{T}_n) F_{n+1} > 0 \quad (3)$$

How can an organization make itself a less attractive target for a serial thief? Most already know that if they have a good relationship with law enforcement, a thief will be more likely to be caught and face conviction. The likelihood of conviction is represented above via \mathcal{P}_c .

By taking account of the actions of past victims in our definition of \mathcal{P}_c , we can also quantify the effect of providing leads to law enforcement and other members of the defense that may not help the victim directly, but may lead to the thief's capture when others are attacked. Thus we can show how, if a victim organization has a reputation for sharing information to help protect others, it can reduce the expected income to the thief from choosing to include it in the set of targets. As a result, organizations that share information with others make less attractive targets than those that keep information to themselves.

Even if leads do not result in the thief's being caught, they may still reduce the thief's expected income from future attacks. The defense may use information learned from one victim to help protect others or to allow the next victim to learn more about how it has been attacked. Such information might include suspicious traffic patterns or a port number the thief is suspected to have used. We see the results of such efforts in our equations as decreases in $\bar{\mathcal{P}}_d(t|\mathbf{T})$ and $\bar{\mathcal{P}}_f(t|\mathbf{T})$ for potential future targets t .

Similarly, having a reputation for detecting exploits used and cutting off any future revenue from using such an exploit once it is patched will also deter future attacks. It would be foolish for an organization to do all this but neglect to try to foil attacks when they happen. As organizations add intrusion detection and response capabilities, thus increasing \mathcal{P}_f , they will also make themselves less attractive targets.

An organization may wish to gauge its attractiveness to a thief in comparison with another organization and pose the question of who would be attacked first. Recall that the later an organization falls in a list of potential victims, the more likely it is that the string of thefts will be brought to an end before the organization is targeted. We can approximate the answer by viewing the world as if there were only two targets, a and b . We write that attacking b followed by a is expected to be more profitable than the reverse ordering by writing:

$$E_{b,a} > E_{a,b}$$

The expected profit from each ordering may be expanded.

$$\begin{aligned} E_{b,a} &= \overline{\mathcal{P}}_f(t_b) \ell_b - \mathcal{P}_c(t_b) F + \overline{\mathcal{P}}_d(t_b) [\overline{\mathcal{P}}_f(t_a|t_b) \ell_b - \mathcal{P}_c(t_a|t_b) F] \\ E_{a,b} &= \overline{\mathcal{P}}_f(t_a) \ell_a - \mathcal{P}_c(t_a) F + \overline{\mathcal{P}}_d(t_a) [\overline{\mathcal{P}}_f(t_b|t_a) \ell_a - \mathcal{P}_c(t_b|t_a) F] \end{aligned}$$

5 Parallel Theft

Thieves undertake parallel attacks in an attempt to penetrate as many systems as possible before the defense can construct a patch for the vulnerability. We assume that the parallel approach ensures that even if the attack is detected, a patch cannot be created and deployed in time to prevent the remaining targets from being penetrated. Hence, the following analysis does not consider \mathcal{P}_d , as defined in the previous section.

Thieves benefit from every target attacked in so far as each target increases the potential loot. For n targets, t_1, t_2, \dots, t_n , we define the marginal loot for the i th target as ℓ_i and refer to the total loot as L_n .

ℓ_i Marginal potential loot from attacking the i th target.
--

$L_n = \sum_{i=1}^n \ell_i$ Total potential loot held by n targets.

As we described earlier, increasing the set of targets also increases the risk of being caught, convicted, and punished. Though detection of a parallel attack may not help the defense at the time of attack, it may prevent the thief from obtaining the loot (for example, if compromised machines are re-secured before the loot is taken) or enable the defense to recover all of the stolen loot at the time of conviction. Thus, a failure or conviction due to attacking any single target will often result in the loss of the loot from all targets. Our analysis incorporates this assumption by setting the total potential loot, L_n , and the total fine paid if caught, F , to be equal.

Given this assumption, we do not want to think about \mathcal{P}_f and \mathcal{P}_c as separate probabilities, but instead simply consider the marginal risk for the i th target as the single quantity r_i . We refer to the probability that the thief successfully captures all of the loot as P_n .

r_i Marginal increase in the probability that all loot is lost due to the attack on the i th target.

$P_n = 1 - \sum_{i=1}^n r_i$ Probability that attack of n targets succeeds and does not lead to detection or conviction.

Using these definitions, we can now state that the expected profit from the attack, E_n , is the potential loot L_n times the probability that the attack succeeds, P_n .

$E_n = P_n \cdot L_n$ Expected profit from the attack.

The expected profit from an additional attack would take into account the marginal loot and marginal risk.

$$E_{n+1} = P_{n+1}L_{n+1} = (P_n - r_{n+1})(L_n + \ell_{n+1})$$

The thief benefits from expanding the set of targets by one if $E_{n+1} > E_n$.

$$\begin{aligned} E_{n+1} &> E_n \\ (P_n - r_{n+1})(L_n + \ell_{n+1}) &> P_n L_n \\ P_n L_n + P_n \ell_{n+1} - r_{n+1} L_n - r_{n+1} \ell_{n+1} &> P_n L_n \\ P_n \ell_{n+1} - r_{n+1} \ell_{n+1} &> r_{n+1} L_n \\ P_{n+1} \ell_{n+1} &> r_{n+1} L_n \end{aligned}$$

The equilibrium point that describes the number of targets attacked, n , is the point at which:

$$P_n \approx \frac{r_n L_{n-1}}{\ell_n} \tag{4}$$

Once the equilibrium point is found, the resulting loot and risk summations, L_n and P_n , can be multiplied to find the expected profit from the attack E_n . Theft is a losing proposition if E_n is less than the cost to find and exploit a new vulnerability. Even if the thief is expected to profit from finding a vulnerability in the system, a given organization may not be targeted if the thief perceives the marginal risk of attacking that organization to be too high.

To help develop intuition for our equilibrium state, consider the case in which all targets contain the same amount of loot ℓ . We can then revise Equation 4 by taking into account that $\ell_n = \ell$ and $L_{n-1} = (n-1)\ell$.

$$\begin{aligned} P_n &\approx \frac{r_n \cdot (n-1)\ell}{\ell} \\ P_n &\approx (n-1)r_n \end{aligned}$$

Dividing by the marginal risk yields the equilibrium state for the case of homogeneous loot:

$$n = \frac{P_n}{r_n} + 1 \quad (5)$$

The conclusion drawn from both the distinct and homogeneous loot cases is the same: maximizing the marginal risk of each additional attack is essential to deterring the parallel thief.

The marginal risk of attack will be low if organizations are defended by homogeneous defenses (packaged systems), as once one system using such a defense is attacked the marginal risk of attacking additional systems that use this defense exclusively will be extremely small. Using a set of systems, each customized to have unique properties, will be much more effective in keeping marginal risk high. This is not surprising as animal species use a diversity of defense mechanisms, customized in each organism, to survive potent threats from parallel attacks.

Sharing of information is also key to keeping marginal risk high. If the body of knowledge of each member of the defense grows with the number of targets attacked, so will the marginal risk of attack. If organizations do not share information, the body of knowledge of each one will be constant and will not affect marginal risk.

As mentioned above, the thief must not only avoid being traced while breaking in, but must also retrieve the loot and exchange it for payment. To keep his marginal risk low, the thief will require both cash and communications systems that provide the maximum level of anonymity available. Anonymous cash fails to protect the thief if the anonymity can be revoked by law enforcement. Anonymous networks may fail to protect a thief transferring large quantities of stolen data, as transfers of this size are likely to be susceptible to traffic analysis. A thief selling back door access to machines or collecting ransoms takes on additional risk with each transaction. Thieves may try to mitigate these risks by finding ways to group transactions.

Once marginal risk has been addressed, the remaining tool in thwarting the parallel attack is limiting the amount of loot lost to the thief. Backing up data both frequently and securely is an excellent way to limit your exposure to both extortion attacks (cryptoviruses), state changing attacks (such as those that manipulate bank balances and try to cover their tracks), and terrorist attacks. Backup systems, considered by most to be a long solved problem, are ripe for a new era of research against this new class of failure.

6 Other threats

We focused on outside thieves because their ability to attack multiple targets makes creating an economic model for their behavior a challenge, but a tractable one. To put this work in a larger context we will briefly mention a few additional classes of adversaries and the types of analyses required.

6.1 Insiders

Insiders have unique knowledge and access that yield a considerable advantage in attacking your organization. It is all but unavoidable that there will be partners, users, administrators, and others in or near your organization that require levels of access that make it easier (cheaper) for them to violate your security policies. This is why insider crime is both so dangerous and so common. However, economic models for insider theft are much simpler than for outsiders, as a typical incident will consist of a single crime against a single target. Whereas one cannot simply compare the amount of the loot available with the *cost to break* of a system when protecting from outside theft, such comparisons can be used when defending against an insider.

6.2 Competitors

What makes competitors unique is that they can dramatically benefit from the target organization's losses. As with insiders, this creates an imbalance in cost benefit ratio of attacking one organization (or a small set of organizations) in comparison to others. Like insiders, a competitor's approach can be more easily analyzed by focusing on a small number of players rather than creating models with an unlimited number of targets. When analyzing how competitors might attack, one should not only look at the vulnerability of each system but also at vulnerabilities introduced through system configuration, system interaction, and your security organization. Because few organizations use the same combinations of interacting systems, it may not be possible to amortize the cost of finding vulnerabilities in these configurations. Thus, the cheapest vulnerabilities to find may lie in configuration, system interaction, or organizational weaknesses. These vulnerabilities are most attractive to competitors, as this class of adversary is not as interested in amortizing the cost of finding a vulnerability as they are in attacking a single target.

6.3 Terrorists

As with outside thieves, an analysis of terrorists must assume that they can and will attack a number of victims, amortizing the costs of finding vulnerabilities and exploiting them. Terrorists are particularly dangerous because, like competitors, they perceive benefit from causing damage regardless of whether they are able to retrieve stolen loot. Owing to differences in motivation and jurisdiction, terrorists are likely to believe they have less to lose if detected or caught.

7 Lessons for the Defense

Our analysis above has focused primarily on strategies for organizations that are targets of the thief. These organizations may be aided by many other members of a larger defense.

7.1 Insurers

Target organizations often benefit from transferring their security risks to insurers. While organizations often do not have the expertise to understand these risks, insurers will require that risks be understood before pricing their policies. Insurers also have the power and incentive to force organizations to pay for better security and to provide the knowledge to help them do so.

Insurers benefit when firms share information about attacks, helping to prevent future attacks from succeeding. To foster this sharing of information, insurance companies may want to offer low deductibles and make payment of claims contingent on timely sharing of information. This not only helps prevent future attacks, but as we saw in the analysis in Section 4 this strategy will make the insured systems less attractive targets for attack by serial thieves.

Parallel-theft attacks may expose insurance companies to many concurrent claims, in the same way that an act of God or terrorist attack would. For this reason insurers may want to think twice before insuring against parallel theft. The risk may be greater if the insurer is partnered with a single monitoring firm, as Lloyd's of London is with Counterpane [3], since a single monitoring firm may provide a more homogeneous defense to the insured assets than would a diverse group of firms.

7.2 Monitoring Firms

Monitoring firms are in a unique position not only to detect known attacks, but to discover the first use of new exploits. Clients of monitoring firms that publish information the moment it is discovered will be less attractive targets than clients of monitoring firms that do not.

These firms also are in a unique position to fight parallel attacks through the use of network traffic analysis. They can use this analysis to detect new viruses and worms, to locate the destination of flows of stolen data, and to detect unusual access requests. This is another ripe area for research, as the ability to thwart worms early in the chain of infection would provide immense value to the firm's clients.

Monitoring firms may also benefit from creating or partnering with Honenets, on whose systems it may be easier to detect parallel attacks. Doing so may improve the monitoring firm's relationship with its insurance partners, or may allow the firm's clients to reduce their insurance premiums.

7.3 Honeynets

Honeypots and Honeynets are closely monitored systems and networks designed by the defense to be infiltrated by the attackers, so that the defense can learn how the attackers is exploiting their systems [4]. Unlike real networks, in which differentiating between an attack and legitimate network use isn't always possible, Honeynets don't have real network traffic. This gives Honeynets a unique advantage in being the first member of the defense to detect a new exploit.

In the name of good citizenship, the creators of Honeynets have used routing rules to keep their compromised systems from being used to attack other machines. This may not be socially optimal. There is no dearth of poorly protected consumer systems available which thieves may turn into 'zombies' used for attack. However, both serial and parallel theft become less profitable with increased risk that a captive machine, used for anonymous routing of stolen goods, may actually be used to detect the exploits (increasing \mathcal{P}_d) or link the thief with a crime (increasing \mathcal{P}_c). Thus, society may instead want to encourage those running closely monitored Honeynets to allow those systems to be compromised and used by the adversary to stage attacks and route data. In particular, conditions under which Honeynets would receive protection from legal liability should be drafted.

Honeynets are currently run by volunteers. For-profit Honeynets may arise to expand the capabilities of the defense if system vendors offer bounties for reports of newly discovered exploits. Honeynets may also prosper through partnerships with monitoring firms and insurers, or by being integrated into these firms.

7.4 Government

Lawmakers ensure that acts of theft and destruction are matched with deterrent punishments, F . Domestic threats should be countered with federal laws and sentencing guidelines and systems for ensuring cooperation between states. This will end reliance on lax state codes, such as those in Massachusetts, where breaking into a system is a misdemeanor with a maximum punishment of \$1,000 and 30 days in prison (Massachusetts General Law, Chapter 266, Section 120F). Once domestic issues are addressed, the true challenge will be to overcome international jurisdictional issues to ensure that $F \neq 0$.

Law enforcement is also essential to fighting theft, especially parallel theft. Whereas serial thieves must contend with the threat that detection will force them to take a loss on their investment in discovering an exploit, the threat of capture after the exploit has done its damage is key to deterring the parallel thief. Once again, jurisdictional issues must be overcome. Lawmakers should discourage the creation of networks and cash systems that have irrevocable anonymity, for if such systems became common the risk of detection in all forms of parallel theft would be greatly reduced. Network traffic monitoring at some level may be necessary to limit the danger of parallel data theft, though such approaches will not be popular with privacy advocates.

8 Related Work

The study of the economics of crime owes a great deal to Becker, who wrote the seminal paper on the subject nearly a quarter century ago [5]. Among his contributions is a model for the expected utility of committing a crime that serves as a foundation for this work. He defines this value by adding to the utility of the crime the product of the utility of the punishment (a negative number) and the probability of conviction.

Ehrlich's work [6] examines both a theory of crime as occupational choice and empirical evidence for this theory. In later work [7] he proposes a market model of criminal offenses, where supply of offenses is a function of the benefits and costs of committing them, including the opportunity cost lost from legitimate labor.

The question of whether it behooves an organization to report information when it is attacked has been posed by Gordon, Loeb, and Lucyshyn [8] in the context of Information Sharing and Analysis Centers (ISACs) such as CERT. A similar question was previously addressed in the context of reporting of household burglaries by Goldberg and Nold [9]. Their empirical analysis found that those households that contained indicators that they would report burglaries were less likely to be robbed. In Section 4 our model suggests that reporting forensic information discovered in response to a systems attack by a serial thief should help to deter future theft.

Anderson [10] has addressed the unfortunate economics of defending against information terrorists in his seminal paper on the economics of information security. Gordon and Loeb [11] examine the optimal defensive response for an individual acting alone, whereas Varian [12] examines optimal behavior of individuals that make up a collective defense in a variety of collaborative contexts.

Detecting vulnerabilities and releasing patches will only protect those who install the patches. Immediate patching is far from a forgone conclusion. Beattie et al. [13] present a formula for system administrators to determine the optimal time to apply a patch. While the formula itself is straightforward, the inputs required (such as the probability of attack, cost of recovery, and potential cost of applying a faulty patch) are likely to be the result of speculation. Rescorla [14] presents a case study showing how slowly the user community reacted to patch a serious vulnerability in Apache both before and after a virus exploiting that vulnerability was released. Much still needs to be done to improve the community's patching process before we can assume that the release of a patch will bring a criminal's spree to a halt.

Finally, this paper relies on the assumption that system security can be measured and that this measurement is the cost to acquire a means of breaking into a system. This method was formally developed by us [1, 15], but owes much to the work of Camp and Wolfram [16], who first proposed that markets for vulnerabilities be created.

9 Conclusion

We have introduced a model for estimating the value of a system exploit to an outside thief. This model takes into account investments in intrusion detection and response, both internally and by outside monitoring firms. Using the model, an organization can gauge its attractiveness to outside thieves and determine how much security is required in the packaged systems it purchases. Beyond choosing how much to spend on security, analysis of the social aspects of intrusion detection and response strategies, such as information sharing, can be evaluated for their effectiveness in deterring future attacks.

10 Acknowledgements

This paper could not have been completed without the advice, comments, and suggestions from Fritz Behr, Glenn Holloway, David Malan, David Molnar, and Omri Traub. This research was supported in part by grants from Compaq, HP, IBM, Intel, and Microsoft.

References

1. Schechter, S.E.: Quantitatively differentiating system security. In: *The First Workshop on Economics and Information Security*. (2002)
2. Young, A., Yung, M.: Cryptovirology: Extortion-based security threats and countermeasures. In: *Proceedings of the IEEE Symposium on Security and Privacy*. (1996) 129–140
3. Counterpane Internet Security, Lloyd's of London: Counterpane Internet Security announces industry's first broad insurance coverage backed by Lloyd's of London for e-commerce and Internet security. <http://www.counterpane.com/pr-lloyds.html> (2000)
4. The HoneyNet Project: Know Your Enemy: Revealing the Security Tools, Tactics, and Motives of the Blackhat Community. Addison-Wesley (2001)
5. Becker, G.S.: Crime and punishment: An economic approach. *The Journal of Political Economy* **76** (1968) 169–217
6. Ehrlich, I.: Participation in illegitimate activities: A theoretical and empirical investigation. *The Journal of Political Economy* **81** (1973) 521–565
7. Ehrlich, I.: Crime, punishment, and the market for offenses. *The Journal of Economic Perspectives* **10** (1996) 43–67
8. Gordon, L.A., Loeb, M.P., Lucyshyn, W.: An economics perspective on the sharing of information related to security breaches: Concepts and empirical evidence. In: *The First Workshop on Economics and Information Security*. (2002)
9. Goldberg, I., Nold, F.C.: Does reporting deter burglars?—an empirical analysis of risk and return in crime. *The Review of Economics and Statistics* **62** (1980) 424–431
10. Anderson, R.J.: Why information security is hard, an economic perspective. In: *17th Annual Computer Security Applications Conference*. (2001)
11. Gordon, L.A., Loeb, M.P.: The economics of information security investment. *ACM Transactions on Information and System Security* **5** (2002) 438–457

12. Varian, H.R.: System reliability and free riding. In: The First Workshop on Economics and Information Security. (2002)
13. Beattie, S., Arnold, S., Cowan, C., Wagle, P., Wright, C.: Timing the application of security patches for optimal uptime. In: Proceedings of LISA '02: 16th Systems Administration Conference. (2002)
14. Rescorla, E.: Security holes... who cares? <http://www.rtfm.com/upgrade.pdf> (2002)
15. Schechter, S.E.: How to buy better testing: Using competition to get the most security and robustness for your dollar. In: Proceedings of the Infrastructure Security Conference. (2002)
16. Camp, L.J., Wolfram, C.: Pricing security. In: Proceedings of the CERT Information Survivability Workshop. (2000) 31–39