Knowing the Enemy at the Gates: Measuring Attacker Motivation

** George P. Corser, Oakland University, United States*
Suzan Arslanturk, Oakland University, United States
Jared Oluoch, Oakland University, United States
Huirong Fu, Oakland University, United States
George E. Corser, Destination Imagination, United States

ABSTRACT

Traditional cost-benefit analysis (CBA) quantifies the value of information security safeguards in terms their expenses compared to their savings before and after their implementation. This paper considers CBA from the attacker's viewpoint, adding another type of measurement, the willingness to endure consequences. We propose a new set of equations and examine their implications vis-à-vis two typical network attacks, identity theft and intellectual property theft.

Keywords

CBA, cost-benefit analysis, copyright infringement, cyber crime, identity theft, information security, intellectual property theft, network security

INTRODUCTION

Sun Tsu said, "Know the enemy, and know yourself, and in a hundred battles you will never be in peril" (Tsu, 1964). What should we know about enemies in battles to cost-effectively protect our applications and data?

Traditional cost-benefit analysis, CBA, would require a review of what data and applications were at risk, and what dollar amounts would be lost if such assets were partially or completely inaccessible, damaged or destroyed. Compare this with the cost of implementing a safeguard. If the safeguard saves more than it costs, then the safeguard should be implemented.

Such analysis requires only that we know the likelihood that an enemy might attack and the price of the assets threatened. Estimates might obtained by statistical evaluation of prior experience, in which case this is essentially self-knowledge. However, shouldn't a quantitative examination of the enemy's motivation also be considered?

The answer is yes, according to our interpretation of Tsu, and that's precisely what this paper is about. We measure an attacker's motivation, commitment, or desire to attack by the minimum dollar amount of suffering he would be willing to endure if there were a 100% probability of getting caught. We call this the breakeven cost of getting caught, or Cost(C). The central thesis of the paper is that attacker motivation, Cost(C), can be quantified, using (1) the known value of a the asset to the attacker, i.e. single gain expectancy, SGE, and (2) the known probability of that the attacker will get caught, Prob(C).

Our model is derived from traditional CBA, which is comprised exclusively of linear equations which consider only monetary costs. Consequently, our model suffers from the same oversimplification as traditional CBA. For clarity of exposition, we disregard all other attack costs and benefits, and we consider examples using only direct assets that have a monetary value. Again, the purpose of the paper is to use equations similar to CBA to better understand the attacker. Our contribution is to present CBA from the attacker’s viewpoint, not to improve or add complexity to CBA.

To our knowledge nothing has been published regarding the formalization of *attacker* benefit within the CBA framework within the field of information technology, IT. Reilly, Rickman & Witt (2012) wrote a fascinating analysis on the benefit of robbing banks. They cited Nobel Prize winner Gary Becker (1974), who wrote seminal work in the field of attacker economic motivation. We believe similar contributions to CBA may aid IT managers in making security decisions.

Why concern ourselves with attacker motivation? Because sometimes it is more cost effective to reduce the value or exposure of an asset than to increase the barricades protecting it.

Consider how we might protect debit cards, which can be used in automatic teller machines (ATMs). Which would be more cost effective, lowering the maximum daily withdrawal limits (MDWLs) or adding dual factor authentication (DFA)? The former reduces the amount of cash a thief might withdraw; the latter makes it harder for a thief to withdraw any amount of cash. Both are safeguards under CBA, but the effectiveness of the former depends directly on the monetary motivation of the attacker, while the latter depends more on the penetration skills of the attacker. That is, the latter makes an attack more difficult, but the former makes it less rewarding.

We assume a rational attacker will not attack at all if the possible penalty (the risk) is too high vis-à-vis the MWL (the reward) regardless of his ability to successfully accomplish the attack. Now the question is: how can we use empirical data to measure attacker motivation?

BACKGROUND AND RELATED WORK

Mercuri (2003), Neubauer, Klemen & Biffl (2005) and Shiau, Hsu, & Wang (2009) demonstrated the utility of CBA in measuring computer related risks and opportunities. But Lee & Shao (2006) showed that sometimes there are drawbacks to using ALE, annualized loss expectancy, a fundamental component of traditional CBA. Neubauer & Hartl (2009) showed, too, that it can be difficult to put CBA into actual practice. So while CBA continues to yield benefits to organizations, there appears to be room for improvement and enhancement.

Conventional CBA starts with a basic question: How much can the organization afford to spend on security? As Whitman (2011) observed, “The approach most commonly considered for a project of information security controls and safeguards is the economic feasibility of implementation."

At least for those systems which protect money or assets quantifiable in dollar terms, security systems are effective to the extent that they save more money than they cost. By adding up prices of hardware, software and personnel we can determine how much money a system costs. But how do we determine how much money a system saves? Below is the traditional cost-benefit analysis (CBA) formula. All equations in this section come from, or are directly derived from, Whitman (2011).

$$CBA = ALE(prior) – ALE(post) – ACS \quad (1)$$

CBA stands for *cost benefit analysis* but it actually means *annualized net benefit*. ALE is annualized loss expectancy. ACS is annualized cost of safeguard. So equation (1) is saying the annualized net benefit (CBA) of a safeguard equals the annualized loss before the safeguard (ALE(prior)), less the annualized loss after the safeguard (ALE(post)), less the annualized cost of the safeguard (ACS). Equation (1) is the fundamental equation of CBA. Suppose you implement a safeguard that costs $5, like a lock. Before the safeguard you lost $100 per year. After the safeguard, you lost zero. CBA = $100 – 0 – 5 = $95. You saved 95 dollars.

Annualized losses are individual losses times the number of total losses. Formally, ALE is the product of SLE, single loss expectancy, and ARO, annualized rate of occurrence.

$$ALE = SLE \times ARO \quad (2)$$

SLE is AV, average asset value, times EF, exposure factor. EF is the percentage of an asset that would be lost if an attack were successful.

$$SLE = AV \times EF \quad (3)$$

Substituting, ALE can be written,

$$ALE = AV \times EF \times ARO \quad (4)$$

This can be restated thus: Annualized loss expectancy (ALE) equals the average asset value (AV) times the percentage of the asset that would likely be lost (EF) times the number of losses in a year.

To put this in more concrete terms, suppose you want to protect a bank's customers' online bank accounts. Assume, before implementing any safeguard, every bank account holds $3000 and every bank account has a $300 per day maximum withdrawal protection. If there are 2000 thefts per year, then ARO=2000. If every account is frozen the day of a theft, then only one $300 theft per account is possible, so EF=$300/$3000=10%. The average asset value is $3000, i.e. AV=$3000. Therefore,

$$ALE(prior) = \$3000 \times 10\% \times 2000 = \$600{,}000 \quad (5)$$

METHOD

ALE(prior) estimates the maximum possible benefit of protecting an asset. If we could achieve 100% protection of an asset, then ALE(post)=0. And if the cost of this protection were zero (ACS=0), then we could say, max(CBA) = ALE(prior) – 0 – 0, or more simply,

$$\max(CBA) = ALE(prior) \quad (6)$$

In the case of our hypothetical bank, if there were a safeguard which could provide 100% protection, it would be worth a maximum of $600,000.

We start from the assumption that inexpensive methods of discouraging attackers provide more cost-effective protection than expensive methods of fortifying defenses.

We propose a set of equations to determine intensity of the motivation of an attacker. These are based on the CBA analysis in the previous section. First, we define SGE, the attacker's single gain expectancy, and CA, the attacker's cost of attacking. SLE remains single loss expectancy, which is the loss suffered by the defender, but also the gain enjoyed by the attacker, less the expenses he incurs by performing the attack. Now we have,

$$SGE = SLE - CA \quad (7)$$

The breakeven point for the attacker is when SGE = 0; that is, when SLE = CA. Going back to our bank example, if it costs an attacker more than $300 to attack one bank account, it's not economically feasible to attack.

Now, let us assume attacks are free except for criminal penalties. Let us ignore any costs of attacking except the penalty for getting caught. In this case, CA can be broken into two parts, Prob(C), the probability of getting caught, and Cost(C), the cost of getting caught.

$$CA = Prob(C) \times Cost(C) \quad (8)$$

Cost(C), or min(Cost(C)), is the cost to the attacker of getting caught, which is a measure of the intensity of motivation of the attacker. Essentially, it can be thought of as the minimum amount of money for which an attacker would willingly suffer the consequences of getting caught.

Suppose an attacker would be willing to spend one year in jail if you gave him $1,000,000, but would not be willing to do so for anything less. Then Cost(C)=$1,000,000. If the chances of an attacker

getting caught are one in a thousand, and the penalty for an attack is one year in jail, then,

$$CA = (1/1000) \times \$1{,}000{,}000 = \$1{,}000 \quad (9)$$

Returning to equation (7), if SLE=$300, the gain expectancy would be negative.

$$SGE = \$300 - \$1{,}000 = -\$700 \quad (10)$$

If the gain expectancy is negative then the attacker would not attack. However, suppose we do not know the attacker's Cost(C). If we know the attacker is aware of the dangers of attacking, and he decides to attack anyway, then we can estimate an upper bound of his Cost(C). In the bank example, breakeven CA=SLE=$300. If Prob(C) remains 1/1000, then

$$Cost(C) \leq CA / Prob(C) = \$300/(1/1000) = \$300{,}000 \quad (11)$$

We may not know exactly what the attacker's minimum acceptable cost would be, but we know that it is less than or equal to $300,000, otherwise he would not have attacked.

CALCULATIONS

We consider two examples, identity theft and intellectual property theft. Using government figures we estimate the motivation level of attackers.

Identity Theft (IDT)

First, consider identity theft (IDT) security measures, first for individuals, then for financial institutions and society.

Identity Theft ALE for Individuals

The Bureau of Justice Statistics (Langton & Planty, 2010) reported 11.7 million people, about 5% of people age 16 or over, were victims of identity theft over a two-year period (Langton, 2010). Assume half, 5.85 million, or 2.5%, were victims during one year. That is, ARO=5.85 million.

We cannot know the exposure factor, EF, or the average asset value, AV. However, we can estimate their product. Financial losses reportedly approached $17.3 billion over two years, about $8.65 billion per year, or $1478.63 per victim. Thus, whatever the asset values were, and whatever the exposure factors were, their product, AV x EF = $1478.63. So from equation (4) we can say,

$$ALE_{IDT}(prior) = AV \times EF \times ARO = \$1478.63 \times 5.85 \text{ million} = \$8.65 \text{ billion} \quad (12)$$

ALE(prior)=max(CBA), so if there were a no-cost safeguard which could eliminate all identity theft then it would be worth no more than $8.65 billion.

Incidentally, only 2.5% of individuals were victims of identity theft. Twenty-three percent (23%) of individuals suffered out-of-pocket costs (Langton & Planty, 2010). Assuming that the 23% who incurred any costs at all in fact incurred the full cost of the thefts, prior to implementing additional security measures, then the portion of the total amount incurred by individuals (not banks) is 23%, and any given individual's estimated dollar cost of identity theft is 23%, so $ALE_{individual}$ (prior) = $1478.63 x 23% x 2.5% = $8.50 per year. We call this CPPPY, the cost per person per year. Suppose security measures were to eliminate 100% of identity theft. In that case, the measures would save $8.50 per individual—for *all* types of identity theft, not just cyber identity theft.

From the average attacker's standpoint, the expected gain is the victim's ALE, not the average person's ALE. On a single loss basis, how can the average person influence the attacker's SGE, which we assume to be equal to the victim's SLE? Monitor credit card transactions more frequently, perhaps. That might limit the number of times an attacker can use the same credit card. But if it takes the defender more than one extra hour per year to keep an eye open for IDT, and his time is worth more than $8.50

per hour, then it's not worth it.

Not all identity theft is computer-related, cyber crime. The BJS reported a wide range of non-computer-related identity theft methods, including 20% of respondents who believed information was stolen from a wallet or checkbook (Langton & Planty, 2010). The United Nations Office on Drugs and Crime published a 2010 report, *The Globalization of Crime*, which estimated the annual volume of identity theft *as a result of cyber crime* at about US$1 billion globally (UNODC, 2010). If the United States represents one-fourth of the world economy, then computer-related identity theft in the US would be $250 million.

Dividing $250 million, cyber IDT, by $8.65 billion, total IDT, yields 2.89%, the percentage of IDT that is cyber IDT. If the proportion of victims is similar to the proportion of money lost, then 2.89% of 5.85 million, or 169,075, would be the expected number of victims of cyber identity theft.

Federal Trade Commission, FTC, publications tend to support this figure. The FTC recorded 250,854 identity theft reports to the Consumer Sentinel Network in 2010, down from 278,356 in 2009 and from 314,521 in 2008 (FTC, 2011). FTC numbers exceed our estimate using UNODC numbers perhaps because the United States, while representing one-fourth of the world's economy, represents greater than one-fourth of the world's cyber-economy.

If 5.85 million victims is 2.5% of the overall population, then the total US banking population is 234 million. If ARO = 169,075 victims divided by a population of 234 million, then expected ARO per banking customer is 0.00072254. If AV x EF remains $1478.63 and 23% of banking customers have out of pocket costs, then expected CPPPY = $1478.63 x 23% x 0.00072254 = $0.25, roughly 25 cents per person for cyber IDT.

If instead, using FTC figures, ARO = 250,854 victims divided by a population of 234 million, then ARO = 0.00107203. If AV x EF remains $1478.63 and 23% of banking customers have out of pocket costs, then CPPPY = $1478.63 x 23% x 0.00107203 = $0.36 per person.

On a per-person basis, cyber identity theft appears to present a small threat. For an average individual, the risk of cyber identity theft by itself may not present sufficiently significant exposure to cost-justify the purchase of additional information security software.

Identity Theft ALE for Financial Institutions

Financial institutions may apply similar cost-benefit calculations. Each Big 4 retail bank, for example, represents roughly 10% of the US retail banking market. Using BJS figures, one-tenth of all victims, ARO=585,000, may be customers of a particular Big 4. Annual costs of identity theft would be 10% of $17.3 billion every two years, or $8.65 billion per year, or AV x EF=$1478.63 per customer victim, the same as for the calculation for individuals.

If banks bear 77% of identity theft costs, then ALE(prior) = $1478.63 x 77% x 585,000 = $666 million, for all banks. While these figures may appear enormous to the lay person, bankers may not consider them worrisome. If the overall banking customer population is 234 million, and each Big 4 bank has 10% of these customers, then each would have 23.4 million customers. Big 4 cost per customer per year, CPCPY, due to all identity theft would be $666 million divided by 23.4 million, or $28.46 per customer per year. The cost due to computer-related identity theft would be $19.25 million divided by 23.4 million, or 82 cents per customer per year, using UNODC calculations; $1.22 using FTC numbers.

Big 4 banks' cost of revenue, COR, averaged $13 billion in 2008 (Yahoo Finance: Bank of America , 2008;Yahoo Finance: JPMorgan , 2008;Yahoo Finance: Citigroup , 2008;Yahoo Finance: Wells Fargo , 2008). Identity theft was a small percentage of this figure. If each Big 4 had 23.4 million customers, then COR was $13 billion / 23.4 million = $555.55 per customer per year. All identity theft (IDT Avg.)

was $28.46 / $555.55 = 5.1% of COR. Computer-related identity theft (Cyber IDT Avg.) was $0.82 / $555.55 = 0.15% (0.0015) of COR, or $1.22 / $555.55 = 0.22% (0.0022) of COR.

What can financial institutions do to reduce the cost of identity theft? Consider one type of identity theft: misuse of credit cards. According to the BJS, misuse of credit cards represented 53% of all identity theft (Langton, 2010). Technical measures do not protect against social engineering attacks or stolen wallets. Most importantly, they do not protect against retailers, online or brick-and-mortar retailers, who do not protect the credit card information of their customers.

Historically, e-commerce companies have presented attractive targets to identity theives. Infamous hacker, Kevin Mitnick, asked, "…what fool would go to all that effort to steal *one* credit card number when many e-commerce companies make the mistake of storing all their customer financial information unencrypted in their databases?" (Mitnick, 2002) Conversely, though security may be lax at certain online retailers, overall risks may be no greater than at storefront operations. Mitnick continued, "There are some hazards to shopping on line, but it's probably as safe as shopping in a bricks-and-mortar store." (Mitnick, 2002) The reason for this, he says, is partly because financial institutions offer their credit card customers protections against fradulent charges.

Using the traditional CBA formula, neither online nor storefront retailers can cost-justify credit card information security. Retailers lose nothing if credit card numbers are stolen. In fact, they profit if identity theives purchase their goods or services. To balance this problem financial institutions may charge higher merchant fees to businesses with higher fraud rates.

Financial institutions have spent on the order of 10% of their annual budgets on information technology, double that of industry in general (Charette, 2005). Additional technical security measures, sufficient to protect 23.4 million customers, are not guaranteed to succeed, and may incur costs in excess of savings to these institutions.

Let's suppose that 100% of stolen credit card numbers are stolen electronically, that none are stolen from wallets. Let us represent as ALE_{CC}(prior) the portion of the annualized loss expectancy due to credit card losses. Then ALE_{CC}(prior) = 0.22% of COR = 0.0022 x $13 billion = $28.6 million. If 53% of this figure is credit-card-related, then ALE_{CC}(prior) = 53% x $28.6 million = $15 million.

Let's further suppose that a system could be created which would protect credit card numbers 100% of the time. In such a case, ALE_{CC}(post) = 0. To break even on a perfect system, CBA = 0. Breakeven ACS, then, would be $15 million. Similarly, if the system only protected 50% of the time, then maximum cost-justifiable ACS would be $7.5 million.

Could a security system be built which would protect 50% of credit card numbers and cost less than $7.5 million to build? If so, financial institutions interests would be served by building such a system. Otherwise, such a system might cost more than it would save.

Identity Theft ALE for Society

Overall, identity theft presents a burden to society commensurate with the sum of the cost to individuals and the cost to institutions. The BJS figure is $17.3 billion over two years, we estimate at $8.65 billion per year. We estimate the cost per person, per member of the banking population, is $8.65 billion divided by 234 million, or $36.97. With UNODC figures, the per-person expected cost is $1.07 per person; using FTC figures, the cost per person is $1.59. See Table 1.

Table 1. Ratio of cyber IDT to total IDT, based on UNODC and FTC numbers.

				Cost Per Person Per Year (CPPPY)		
Source	Cyber ARO	Total ARO	Ratio	Individual	Bank	Society
UNODC	169,075	5,850,000	2.89%	$ 0.25	$ 0.82	$ 1.07
FTC	250,854	5,850,000	4.29%	$ 0.36	$ 1.22	$ 1.59

Attacker Calculations, IDT

In the case of IDT, losses to the defender mean gains to the attacker, or SLE = SGE = $1478.63. One-out-of-700 has been estimated as the probability of getting caught for committing identity theft (Litan, 2003). From these numbers we can estimate from equation (8) the breakeven point for attackers,

$$\$1478.63 = (1/700) \times \text{Cost(C)} \tag{13}$$

Therefore, from the inequality in (11) we can say,

$$\text{Cost(C)} \leq \$1{,}035{,}041 \tag{14}$$

Using this reasoning, from an attacker's perspective, if he would be willing to suffer the consequences of getting caught for less than $1,035,041 then he should attack.

Intellectual Property Theft (IPT)

Intellectual property theft covers the federal civil IP cases that can be classified into copyright, patent and trademarks according to the Bureau of Justice Statistics report. Since the cyber intellectual property theft cases are our main concern in this paper, we will only study cost benefit analysis of copyright infringement.

Copyright infringement involves counterfeiting or pirating digital content such as digital music, digital movies and software. The copyright infringement occurs when a copy of a song, movie or any type of software on a CD is being shared on internet using any file sharing network, so people all around the world can download it.

IP Theft Losses, Dollar Estimates

Siwek (2007), in an International Federation of the Phonographic Industry, IFPI, study reported that the music sales diminished by almost $15 billion between 1999 and 2008, with the increase in file sharing networks. The IFPI report showed that recorded music digital piracy is between $17 billion - $40 billion in 2008. IFPI also reported that 40 billion files per year are being downloaded. This information is found by finding the number of people that downloaded illegal music and the number of music files they downloaded per month. Siwek's analysis suggests that the per-incident dollar cost of copyright infringement is $1 per incident, i.e., SLE = AV x EF = $1.

Movie revenues did not suffer as much as the music revenues. Again, according to the IFPI, movie piracy cost between $10 billion - $16 billion in 2005. The software industry suffered from the digital piracy as well. The IFPI stated that software piracy was around $1.5 billion - $19 billion. (Siwek, 2007)

The Recording Industry Association of America, RIAA, quotes the IFPI and other sources on its website (RIAA, 2013), but does not endorse a specific estimate. The closest thing we have to an estimate from the RIAA is, "… U.S. Internet users annually consume between $7 and $20 billion worth of digitally pirated recorded music." At $1 per download, that's 20 billion downloads (ARO=20 billion). While this estimate is rough, we will use it, as we intend to show that even dramatically varying the ARO figure yields a very high Cost(C).

IP Theft ALE, Music Only

Since the profit loss due to digital music piracy itself represents a huge proportion of the total profit loss of cyber intellectual property theft, we consider music piracy only. From the numbers above,

$$ALE_{IPT\text{-}Music}(Prior) = \$20 \text{ billion/year} = \$1 \text{ x } 20 \text{ billion/year} \quad (15)$$

Dividing this number by the US population of roughly 300 million yields the average annual music theft cost of per person: $66.67.

Attacker Calculations, IPT

From the attacker's point of view, consequences of getting caught must be estimated. The Bureau of Justice Statistics report (Motivans, 2004) states that the jurisdiction of civil intellectual property complaints filled in the U.S. district courts for the copyright infringement was 2084 private cases. Since we estimate the number of downloads for music piracy is around 20 billion, the probability of being caught, Prob(C), is 2084/20,000,000,000, or around 0.0000001042. The 2084 private cases cover the music, movie and software industries. So the probability of being caught for just downloading illegal music is even less than 0.0000001042.

We can now estimate Cost(C) from (11).

$$Cost(C) \leq \$1 / 0.0000001042 = \$9{,}596{,}929 \quad (16)$$

So if an attacker would be willing to endure the consequences of getting caught for a sum less than or equal to roughly $10 million, then he should go ahead and download the music. If the number of downloads is in reality half our estimate, then Cost(C) may only be $5 million; if double, then perhaps Cost(C) is on the order of $20 million.

Four Scenarios, using Known Cost(C) to find Prob(C)

Another way to examine the enemy is to find the breakeven Prob(C). Consider some defensive scenarios using different values for Cost(C).

If SGE=$1 and Cost(C)=$200, then Prob(C)=0.005, one out of 200. In other words, if there were a 0.005 chance or less of getting caught then the attacker should attack, otherwise he should not.

If SGE=$1 and Cost(C)=$150,000, then Prob(C)=0.0000066667, one out of 150,000. In other words, if there were a 0.0000066667 chance or less of getting caught then the attacker should attack, otherwise he should not.

If there are $20 billion worth of illegal music downloads, and each music download is worth $1 then there are 20 billion illegal downloads.

To achieve a 0.005 probability of getting caught, the music industry would need to prosecute 0.005x20 billion incidents, or 100,000,000 cases. If each case required the expenditure of $1000 to successfully prosecute, then it would cost $100 billion dollars in legal fees, five times more than the loss being suffered. On the other hand, if the music industry actually collected on every successful prosecution, they would win $200 times the number of successful cases, $200 x 100,000,000. That's $20 billion. The result would be an $80 billion net loss. Under such a scenario, prosecution of offenders would not be a cost effective security measure. (See Table 2, Scenario 1)

To achieve a 0.0000066667 probability of getting caught, the music industry would need to prosecute 0.0000066667x20 billion incidents, or 133,333 cases. If each case required the expenditure of $1000 to successfully prosecute, then it would cost $133,333,333 dollars in legal fees, 1/150th of the $20 billion loss being suffered. If the music industry actually collected on every successful prosecution, they would win $150,000 times the number of successful cases, $150,000 x 133,333. That's $20 billion. The result would be a $19,866,666,667 net gain. Under such a scenario, prosecution of offenders would be a cost effective security measure. (See Table 2, Scenario 2)

The analyses above assume that every case is prosecuted successfully, with no lost cases. Could the recording industry prosecute 133,333 cases of music piracy and win $150,000 every time, with no lost cases, and each case cost only $1000 in legal fees to prosecute? If so, then they should prosecute.

Table 2. Cost-benefit of prosecuting digital music theft: four scenarios

Cost-benefit item	Scenario 1	Scenario 2	Scenario 3	Scenario 4
SLE: Dollars lost per incident	$ 1	$ 1	$ 1	$ 1
ARO: Incidents per year	20,000,000,000	20,000,000,000	2,000,000,000	2,000,000,000
ALE(prior): Total dollars lost per year	$ 20,000,000,000	$ 20,000,000,000	$ 2,000,000,000	$ 2,000,000,000
Cost(C) *	$ 200	$ 150,000	$ 200	$ 150,000
Breakeven Prob(C)	0.005	0.0000066667	0.005	0.0000066667
Cases to prosecute to achieve Prob(C)	100,000,000	133,333	10,000,000	13,333
Cost per successful prosecution **	$ 1,000	$ 1,000	$ 1,000	$ 1,000
ACS: Legal fees = Cases x Cost per case	$ 100,000,000,000	$ 133,333,333	$ 10,000,000,000	$ 13,333,333
Winnings per successful prosecution	$ 200	$ 150,000	$ 200	$ 150,000
Total winnings = Cases x Winnings per case	$ 20,000,000,000	$ 20,000,000,000	$ 2,000,000,000	$ 2,000,000,000
ALE(post) = ALE(prior) - Total winnings	$ 0	$ 0	$ 0	$ 0
CBA = ALE(prior) – ALE(post) – ACS	$ (80,000,000,000)	$ 19,866,666,667	$ (8,000,000,000)	$ 1,986,666,667

* assume Cost(C) is known, ** assume every case prosecuted is successful (no lost cases)

These analyses show, not the Cost(C) threshold, but the Prob(C) threshold. That is, if we know the cost of getting caught to the perpetrator is worth $200 then we would have to prosecute and win 100,000,000 cases out of 20 billion to secure the digital music. If the cost of getting caught to the perpetrator is worth $150,000 then we would have to prosecute and win 133,000 cases out of 20 billion to secure the digital music. In other words, both the Cost(C) computation, $9,596,929, and the Prob(C) computations, 0.005 or 0.0000066667, suggest that prosecution is not a cost effective method of securing digital music.

Calculations using Nonmonetary Units

Like traditional CBA, the attacker CBA model presented in this paper is limited to monetary benefits and costs, which are not always the critical factors from an attacker's viewpoint. Consider the example of web defacement, or vandalism in general for that matter. Let's say the probability of getting caught is one in 700 and the dollar value of the asset is zero. Then Cost(C)≤0/(1/700)=0. That is, the attacker must have been willing to suffer the consequences of getting caught for no monetary benefit in return.

This is intuitively obvious, but there are several ways to interpret this figure. One is to assume all such attackers are irrational. Another is to assume the attacker-CBA model is just plain wrong. Still another is to assume that such attackers are simply not financially motivated; that defenses attempting to reduce the dollar value of the target to the attacker would likely be ineffective. Let's consider this latter possibility, that the attacker is rational but not money-motivated.

The model we are proposing is not designed to apply under circumstances where costs and benefits cannot be measured in money terms. But suppose, instead of monetary currency we used a different "currency," such as pride or revenge. We do not know how to define pride-units or revenge-units, but suppose we could. Further suppose getting caught could be measured in pride-units and defacing the website could be measured in pride-units. If defacing the site would yield to the attacker 1000 pride-units, then Cost(C)≤1000/(1/700) = 700,000 pride-units. The attacker should attack if he would willingly suffer the consequences in exchange for 700,000 pride-units or more. The model still works, if the units change.

The non-monetary-units problem is a serious deficiency, of course. However it is not unique to our model. Neither CBA nor attacker-CBA applies in situations where assets have non-monetary values. Both models measure exclusively using monetary costs and benefits.

CONCLUSIONS

"Know the enemy, and know yourself, and in a hundred battles you will never be in peril" (Tsu, 1964). Traditional CBA helps IT managers estimate the value of assets and the likelihood of attacks, then compare safeguard costs against expected losses, that is, to "know yourself." We have shown that it is sometimes also useful to apply similar calculations from the attacker's viewpoint, i.e. to "know the enemy." Identity theft and music theft are both crimes, but their differing attacker cost profiles suggest different attacker profiles, and perhaps different defensive measures.

In our first example, the cost of identity theft averages $1478.63 per incident (AV = $1478.63). If the probability of getting caught is 1 in 700, and the cost of getting caught averages one year in jail, then an attacker should be willing to go to jail for one year for $1,035,041, otherwise he should not attack.

These numbers represent very rough estimates, of course, and future work will hopefully more accurately estimate such costs. But based on these rough estimates, compared to digital music thieves, identity thieves appear to be more willing to suffer consequences such as jail time. Another way of putting it is there is a narrower range of people willing to attack credit cards because their willingness to suffer consequences must be lower.

What might this mean to defenders? Increasing monetary penalties and increasing likelihood of getting caught should always reduce the number of attacks. This must be weighed against the cost of prosecution. Prosecution is only cost effective if defendants have money. If defendants do not have money, then barricades may be more cost effective than legal consequences. Cyber identity theft is especially difficult to prevent if users have easy-to-guess passwords (Mitnick, 2006). Dual factor authentication may be indicated, but this may inconvenience bank customers. Some contend that the most cost effective solution is to write off identity theft loss as a cost of doing business (Newman, 2005). They argue the cost of perfect identity theft defenses does not make financial sense.

Future research must clarify attacker profiles. Are attackers generally family members who have physical access to credit cards, or know the answers to security questions that enable them to bypass password protections? Or are attackers usually anonymous, international, hard-to-prosecute organized gangs? Benefits of defensive measures are limited by our understanding of the attackers.

In our second example, the cost of music theft is roughly $1 per song (AV = $1). If the probability of getting caught and punished is 0.0000001042, and the cost of getting caught averages one year in jail, then an attacker should be willing to go to jail for one year for $9,596,929, otherwise he should not attack.

Though the dollar value of each individual theft is lower, the risk/reward balance of music theft appears better for attackers than the risk/reward balance for identity theft. This indicates a broader range of attackers.

Table 2 shows prosecution may not be a cost-effective means of protecting digital music. Considering the sheer number of songs illegally downloaded, future study may be indicated in the field of economics. Perhaps volume pricing for digital music consumers might capture more revenue for music producers. The broad range of attackers indicated by the cost profile, coupled with the ineffectiveness of prosecution, suggests that at least some of the thefts are perpetrated by people who also purchase songs legitimately. Just as people who like to chat on mobile phones purchase plan with more minutes, perhaps digital music providers could reduce the prices of songs for high volume listeners.

Again, understanding the attacker, knowing the enemy, might improve the cost effectiveness of defenses. Are digital music pirates bootlegging copies and reselling for personal gain, in cash? Or are they sharing music with friends to avoid running up a huge iTunes bill?

In sum, traditional CBA helps IT managers understand themselves. We hope our model for attacker CBA helps IT managers better understand the enemy.

REFERENCES

Becker, G. S. (1974). Crime and punishment: An economic approach. In Essays in the Economics of Crime and Punishment (pp. 1-54). UMI.

Charette, R. (2005, September). Why software fails. IEEE Spectrum, Retrieved from http://spectrum.ieee.org/computing/software/why-software-fails/0

FTC. Federal Trade Commission, (2011). Consumer sentinel network data book for January - december 2010. Retrieved from website: http://www.ftc.gov/sentinel/reports/sentinel-annual-reports/sentinel-cy2010.pdf

Gartner. (2003, July 21). Gartner says identity theft is up nearly 80 percent; 7 million u.s. adults were identity theft victims in the past 12 months. Retrieved from http://www.businesswire.com/news/home/20030721005204/en/Gartner-Identity-Theft-80-Percent-7-Million.

Internet Crime Complaint Center (IC3). U.S. Department of Justice, Bureau of Justice Assistance. (2010. 2008 internet crime report. Retrieved from website: http://www.ic3.gov/media/annualreport/2008_IC3Report.pdf

Langton, L., & Planty, M. U.S. Department of Justice, Office of Justice Programs. (2010). Victims of identity theft, 2008 (NCJ 231680). Retrieved from website: (Whitman, 2003) http://bjs.ojp.usdoj.gov/index.cfm?ty=pbdetail&iid=2222

Lee, V. C., & Shao, L. (2006). Estimating potential IT security losses: An alternative quantitative approach. Security & Privacy, IEEE, 4(6), 44-52.

Litan, A. (2003, July). Underreporting of Identity Theft Rewards the Thieves, Stamford, CT: Gartner Research http://www.gartner.com/press_gartner/images/116066.pdf.

Mercuri, R. T. (2003). Security watch-analyzing security costs. Communications of the ACM-Association for Computing Machinery-CACM,46(6), 15-18.

Mitnick, K., & Simon, W. (2002). The art of deception. Hoboken, NJ: John Wiley & Sons.

Mitnick, K., & Simon, W. (2006). The art of intrusion. Hoboken, NJ: John Wiley & Sons.

Motivans, M. (2004). Intellectual property theft, 2002. US Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

Neubauer, T., Klemen, M., & Biffl, S. (2005, May). Business process-based valuation of IT-security. In ACM SIGSOFT Software Engineering Notes (Vol. 30, No. 4, pp. 1-5). ACM.

Neubauer, T., & Hartl, C. (2009, June). On the singularity of valuating IT security investments. In Computer and Information Science, 2009. ICIS 2009. Eighth IEEE/ACIS International Conference on (pp. 549-556). IEEE.

Newman, G. R., & McNally, M. M. (2005). Identity theft literature review. United States Department of

Justice: National Institute of Justice.

RIAA. (2013). Riaa - anti-piracy. Retrieved from http://www.riaa.com/faq.php.

Rutledge, R., Lalor, A., Oller, D., Hansen, A., Thomason, M., Meredith, W., Foil, M., & Baker, C. (1993). The cost of not wearing seat belts; a comparison of outcome in 3396 patients. Annals of Surgery, Retrieved from http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1242750/

Shiau, W. L., Hsu, P. Y., & Wang, J. Z. (2009). Development of measures to assess the ERP adoption of small and medium enterprises. Journal of Enterprise Information Management, 22(1/2), 99-118.

Siwek, S. E. (2007). The true cost of sound recording piracy to the US economy. Institute for Policy Innovation.

Tsu, S. , & Griffith, S. (1963). The art of war: Translation by Samuel B. Griffith. Oxford, UK: Oxford University Press.

U.S. Census Bureau, (2012). 2012 statistical abstract, transportation: Motor vehicle accidents and fatalities. Retrieved from website: http://www.census.gov/compendia/statab/2012/tables/12s1103.pdf

UNODC. United Nations, Office on Drugs and Crime. (2010). The globalization of crime. Retrieved from website: http://www.unodc.org/documents/data-and-analysis/tocta/TOCTA_Report_2010_low_res.pdf

Whitman, M. (2003). Enemy at the gates: Threats to information security. Communications of the ACM,46(8), 91-95.

Whitman, M., & Mattord, H. (2010, October). The enemy is still at the gates: Threats to information security revisited. Paper presented at InfosecCD '10, New York, NY

Whitman, M., & Mattord, H. (2011). Principles of information security. (4 ed.). Stamford, CT: Cenage Learning.

Yahoo Finance. Bank of America Income Statement. 2008. http://finance.yahoo.com/q/is?s=BAC&annual.

Yahoo Finance. Citigroup Income Statement. 2008. http://finance.yahoo.com/q/is?s=C&annual.

Yahoo Finance. JPMorgan Chase Income Statement. 2008. http://finance.yahoo.com/q/is?s=JPM&annual.

Yahoo Finance. Wells Fargo Income Statement. 2008. http://finance.yahoo.com/q/is?s=WFC&annual.