Three Case Studies in Quantitative Information Risk Analysis

Mohammed A. Bashir and Nicolas Christin
Carnegie Mellon University, INI/CyLab Japan
ashbashir@cmu.edu, nicolasc@cmu.edu

In this paper, we build on existing literature and on a dialog with several decision-making partners (e.g., CISOs) to propose a simple methodology to quantitatively assess the value of security. We use this methodology to provide quantitative data gathered from three case studies of real organizations. The vastly different results we obtain across the three organizations considered emphasize the dependence between the security investments and the nature of the organization implementing them.

1 Introduction

Implementing security is potentially costly, may be partially ineffective, and does not generate any direct revenue for an organization. In addition, organizations are faced with trade-offs when they consider mitigation strategies to prevent attacks. A countermeasure may mitigate an attack, but is also likely to make tasks for the organization's end users more difficult. Take the example of spam: Not only it is extremely difficult for a spam filter to block all undesirable emails, but the spam filter may also block legitimate traffic, further impeding productivity.

As such, convincing non-technical decision-makers to invest in security is a daunting task for the technical managers or security officers who have to justify security expenditures. It is nevertheless a mandatory undertaking to avoid monetary losses due to security breaches [10].

Peltier [9] argues that, in information security, qualitative risk analysis is far easier to conduct than quantitative risk analysis, notably due to the complexity of the computations involved in quantitative models, and the lesser amount of security expertise needed. A further criticism against quantitative models is that, while offering seemingly precise estimates of the damage and recovery costs, they more often than not have considerable margins of error, due to the core assumptions on which they rely.

However, the key advantage of quantitative models is that they provide an actual dollar amount or "bottomline," which makes them appealing to non-technical decision makers. Among quantitative models designed for information security risk analysis, we can cite (in chronological order) the models proposed by Meritt [8], Tan [12], Blakley [4], Greer et al. [6], or Arora et al. [3]; but it is worth noting that most organizations doing quantitative risk analysis use their own model, tailored to their own specifics [9]. Publicly available decision-aid tools, for their part, have been either focusing on qualitative aspects [2], or, on the other hand, on very specific aspects, e.g., network topology [11], consistency between risk analysis and investments [1].

This paper argues in favor of quantitative models for information security. We rely on a simple methodology, described in Section 2, and present three case studies based on actual organizations in Section 3. Our case studies outline the dependencies between the size of an organization, the threats it faces, and the security measures it has in place. We discuss our results and conclude in Section 4.

2 Methodology

To describe our case studies, we rely on a simple quantitative model. A key feature of the model is its simplicity, motivated by usability constraints: we want it to be available as a tool usable by technical as well as management personnel. As such, our model relies on a few numbers to input, and provides a relatively small set of output metrics. Intermediate calculations may be relatively complicated, but can be automated. An implementation of our model is available as an Excel spreadsheet from `http://arima.okoze.net/isra`.

As simple as it may be, our model tries to address a wide variety of threats, not just IT risks. For instance, it also tries to capture risks posed to information held in different media (e.g., paper), and takes into account a comprehensive list of threats ranging from corrupt backup to dumpster diving.

This model is the product of an iterative approach: we created an original version of the model, using, as a basis, the information security risk analysis framework developed at LBNL [3]. We then presented our model to ten external partners, whose backgrounds and affiliations cover a fairly large range, both from the technical and management aspects; partners include researchers, security managers, and well as senior management, and are located in Asia, Europe, United States, and the Middle East. After gathering and integrating feedback from our partners, we revised our model, and arrived at the version we describe next.

The model revolves around attacks, outcomes of these attacks, countermeasures, countermeasure efficacy, and uses two primary types of output: (a) a Value at Risk (VaR) analysis, based on differing countermeasures, and (b) a Risk-based Return on Investment (RROI), that is, the ratio between the net benefit in implementing countermeasures and the cost for such countermeasures) of security controls [3].

The rationale for the Value at Risk analysis is that it is seemingly the best understood language from the financial community. At the same time, Risk-based ROI measures how effectively an organization uses its resources to avoid or reduce risk, and appears a necessary input for budgeting considerations.

Attacks and countermeasures To arrive at the VaR and RROI, we first take the attack types as input and the frequency of attacks over the past year. Using this and the percent coverage of recorded data, we calculate the estimated number of attacks per year.

Here our model uses a key assumption, that the recent past is a good indicator of what will happen in the near future. In other words, the threats are not expected to change drastically from one year to the next. While this assumption may be considered relatively stringent, it is relatively hard to avoid it without resorting to pure speculation about future events.

For attack i, consider the attack frequency F_i, and the coverage $G_i \leq 1$. The coverage corresponds to an assessment of how much data has been recorded, compared to the number of incidents that actually happened; ideally G_i should be equal to 1, but we need to take into account possible weaknesses in the audit trail maintenance. The estimated number of attacks per year is $A_i = F_i/G_i$.

The model also takes the countermeasures in place as input along with the effectiveness of each of the countermeasures against each of the attack types defined above (denoted by C_{ij} for countermeasure j against attack i) . Using this data and the attack frequency, we can calculate the estimated number of attacks that the organization would have suffered had there not been any countermeasures in place.

First, the probability all countermeasures fail against an attack i is $CF_i = \prod_j (1 - C_{ij})$. This formula makes the assumption that countermeasures are independent of each other. This assumption is reasonable for countermeasures that are largely orthogonal, for instance, physical security on the one hand, and data encryption on the other hand.

From there we get the number of attacks we would have in absence of countermeasures, $B_i = A_i/CF_i$, and finally the number of attacks prevented by countermeasures $R_i = B_i - A_i$.

Attacks vs. attack outcomes There is a crucial difference between an instance of an attack (e.g., malicious code infection), and its outcome (e.g., unavailability of a user PC). Countermeasures can thwart attacks; but, only attack outcomes affect the organization's bottomline. The relationship between attacks and attack outcomes is given by a matrix (α_{ij}). For instance, for attack i denoting a malicious code infection, there may be two possible outcomes: destruction of all information with a probability of 100%, and unavailability of the user PC with a probability of 70%. Then, we have $\alpha_{i,1} = 1$, $\alpha_{i,2} = 0.7$, and $\alpha_{i,j} = 0$ for any different outcome j (e.g., unavailability of a print server).

Losses We now shift to the expected losses. Consider the annual salary of an IT employee, M, the employee cost per day is estimated to be $P = 1.5M/365$, where the 1.5 factor has been chosen after discussion with partners to take into account tax and administrative overhead. This cost P will lead us to the expect loss per attack, once we have estimated the number of workdays an attack costs.

We use, as an input, the number of one-person days it would take, for an IT professional, to perform the following type of efforts: 1) The effort needed to diagnose a typical attack, 2) The effort needed to report a typical attack, 3) The effort needed to repair the damage caused by a typical attack, and 4) The effort needed to address any public relations/reputation issues arising from a typical attack.

Here, a typical attack refers to an attack that is most common, that is, one that is closest to the median with respect to severity. With this information, for each attack outcome j, we get the nominal damage N_j as

$$N_j = \sum_k D_{jk} P + C \ ,$$

where C is a parametrized cost noted for other, collateral attack damage not taken into account in other calculations, and D_{jk} represents the attack damage (in days), with k representing the four types of efforts (diagnosis, report, repair, follow-up).

We go from the nominal damage to the expected loss per attack outcome, EL_j by considering the extra severity $S_{j,*}$ of the attack. This is done by specifying the probability that any given attack outcome will be ten ($S_{j,10}$), one hundred ($S_{j,100}$) or one thousand ($S_{j,1000}$) times more severe than the typical outcome(s) for an attack of that type. This accounts for attacks that sometimes result in a high degree of damage. We get

$$EL_j = N_j[10S_{j,10} + 100S_{j,100} + 1000S_{j,1000} + (1 - S_{j,10} - S_{j,100} - S_{j,1000})] \ ,$$

as the expected loss for attack outcome j. This metric is equivalent to the Annual Loss Expectancy for attack outcome j.

Combining the previous outputs, we can calculate the expected loss without countermeasures and the loss avoided due to the use of countermeasures: With these two pieces of information it is now possible to calculate the residual loss per attack outcome (ELC_j). This is the same as the expected loss with countermeasures in place. We have

$$ELC_j = EL_j \sum_i \alpha_{ij} A_i \ \ ,$$

for a total residual risk $RR = \sum_j ELC_j$.

The sum of the expected loss for each attack type yields the total estimated expected loss that the organization incurred in the previous year. This can also be considered the

total residual risk the company is exposed to. All factors being the same, the organization is likely to incur this cost or loss from the attacks specified over the next year.

The expected loss per attack outcome (without countermeasures) is, likewise, $ELwoC_j = EL_j \sum_j \alpha_{i,j} B_i$, and the loss avoided thanks to the countermeasures is simply $LA_j = EL_j - ELwoC_j$.

Benefit of countermeasures We have so far looked at residual risks, but have not assessed the benefit associated with a given type of countermeasure. Consider the cost of capital r, and a time period in years given as t. Consider the total expected loss without any countermeasure, $ELwoC = \sum_j ELwoC_j$, then the benefit associated with only countermeasure k being in place is $BC_k = ELwoC - LC_k$, where LC_k is the total loss when only countermeasure k is in place. From BC_k we can get the current NPV for countermeasure k, NPV_k, as

$$NPV_k = BC_k - CM_k(1-r) \ ,$$

and the NPV over r and t as

$$NPV_{k,r,t} = \sum_l \frac{BC_k - CM_{k,l}}{(1+r)^t} - CM_k(1-r) \ ,$$

where $CM_{k,l}$ is the ongoing cost of countermeasure k, over interval l.

We can also compute the residual risk for each countermeasure acting alone. This is combined with the cost of the countermeasure to produce the net benefit of the countermeasure, and then the ROI for the countermeasure. The net benefit for countermeasure k is $NBCM_k = BC_k - CM_k$, which gives use a ROI for countermeasure k of $ROIC_k = NBC_k / CM_k$.

Simulating the value at risk Some of the inputs, in particular those pertaining to attack severity and number of occurrences, may be subjective. As such, the residual risk computed may be inexact. To solidify the predicted values, we complement the residual risk calculations with Monte-Carlo simulations of the value at risk.

We take the estimated number of attacks and the probability that this figure is 50% higher and 50% lower than the estimate. This data is used as input into a binomial distribution function with a random value to calculate a new attack frequency. This new attack frequency is then input into the model and a new total residual risk value is calculated. This procedure is repeated for a large number of n instances. This then allows for a Value at risk calculation for a specified confidence level.

3 Case studies

We next turn to a description of three case studies on which we use our model to gain a better understanding of the intricacies between each situation (security threats, particularities of the organization), and the effectiveness of selected countermeasures. The first case study is of a small network solutions company, the second case study is of a non-profit organization in the UK, and the third case study is of a major project within a Japanese insurance company. The full input and output stages of the model are available as an online appendix at `http://www.andrew.cmu.edu/user/nicolasc/publications/isra-appendix.pdf`.

3.1 Small network solutions company

This case study is for a small network solutions company. The company has an annual turnover of $4.8m and 22 employees. A typical IT employees salary is $31,000, which

Table 1: **Case study 1: Countermeasure effectiveness**

Countermeasure	Cost	ROI	Curr. NPV	NPV w/ r, t
Anti-virus	$1,000	1057%	$10,728	$8,349
Firewall	$2,000	-22%	-$135	-$339
IDS	$600	219%	$1,403	$1,154
Training and education	$1,500	1437%	$21,777	$18,770
UPS	$2,500	-78%	-$1,571	-$1,643
Active directory	$1,000	956%	$9,709	$8,332
Backup server	$1,200	-51%	-$435	-$511
Spam filtering	$500	1179%	$5,969	$5,135
Network access ctrl.	$2,200	398%	$9,083	$7,654
Email policy enforct.	$2,000	223%	$4,758	$3,916

Table 2: **Case study 1: Expected loss per attack outcome**

Attack Outcomes	EL per Attack Outcome
Information Theft/Disclosure	$322.93
Information Modification	$1,145.85
Information Destruction	$1,178.58
Service (User PC) Unavailable	$211.19
Legal/compliance problems	$6.46

equates to an IT employee cost per day of $129. The company faces the following attacks/threats: (1) Malicious code infections, (2) Administrator account compromise, (3) Regular account compromise, (4) Improper use, (5) Theft, (6) Spam, and (7) Natural disaster.

We calculate the residual risk to be $33,819. This is the amount that the company can expect to lose through the attacks we have considered in our model over the next year, assuming that the attack frequencies, attack outcomes, attack-attack outcome relationships, and countermeasure effectiveness remain the same.

With 95% confidence, we can infer that over the next year the residual risk (the amount the company is likely to lose) will be no more than $41,968. And with 99% confidence we can determine that over the next year the residual risk will be no more than $46,107.

Armed with the Value at Risk, senior management can now decide whether they wish to accept the risk or attempt to reduce it. If they attempt to reduce it, we can again use the model to estimate the ROI and NPV for additional countermeasures.

Table 1 shows the ROI, current NPV and NPV, where $r = 0.15$ and $t = 3$ years for each of the countermeasures in place acting alone. The ROI and NPV for most of the countermeasures is positive, showing these countermeasures are cost-effective. However, the firewall, backup server, and UPS seem not to provide good value for the services they provide.

The ROI for the countermeasures calculated is the ROI for each countermeasure acting alone. Our simple model does not take into account interactions between countermeasures, which may be particularly complex. They are dependent upon the combinations of countermeasures and the network architecture and configurations. Our model assumes that the combined effectiveness of the countermeasures is multiplicative, which may be overly optimistic.

Table 2 identifies the attack outcomes that result in the highest expected loss per attack outcome. The attack outcome with the highest expected loss is information destruction, and is therefore what the manager should try to prevent as much as possible. We can now identify the attacks that lead to the attack outcome, and consider countermeasures that will

Table 3: **Case study 1: Attacks that lead to information destruction and their respective losses**

Attack	Freq.	% result. in info. destruction	Expect. loss
Malicious code infection	20	35%	$8,250
Improper use	30	10%	$4,420
Natural disaster	2	10%	$236

help mitigate these attacks, thus reducing expected loss and residual risk.

Table 3 shows that the attack that results in the highest expected loss for an information destruction attack outcome is the malicious code infection. Using this information, we can explore possible countermeasures to mitigate against malicious code infections. The company currently has an anti-virus in place that is 90% effective in mitigating malicious code infection attacks. The company could invest in a more advanced secondary screening process for files that enter the system; essentially a second anti-virus.[1]

Assuming that we have a new secondary anti-virus (AV2, costing $3,000), which the vendor claims will be 80% effective against the malicious code infections that the company currently faces, the number of malicious code infection type attacks will be reduced from 20 to 4. This leads the residual risk to become $23,527 given the addition of the new anti-virus. Therefore, the benefit from the new anti-virus is $10,292 ($33,819 - $23,527). The net benefit is $7,292 ($10,292 - $3,000). We can then also calculate the ROI for AV2 as follows:

$$\text{ROI for AV2} = \frac{\text{prev. RR - new RR - cost of AV2}}{\text{Cost of AV2}} \approx 242\% \ .$$

The NPV is

$$\text{NPV} = \frac{\text{prev. RR - new RR - ong. cost of AV2}}{(1+r)^t} - \text{cost of AV2} \ .$$

With a capital cost of 15%, a time period of 2 years and an annual cost of $500, the NPV is roughly $4,404.

Using the same method the company can calculate the ROI and NPV of another anti-virus solution, and see which of the two is better.

Another point to note is that the new profit expected from ventures that have been profitably undertaken, thanks to the countermeasure, are not taken into account in the ROI and NPV calculations. These are projects that would not have been possible due to an excessively high risk exposure had the countermeasure not been in place. This is the opinion espoused by Soo Hoo [7], who calculates ROI simply as the annual benefit over the cost of the countermeasure. Blakley [4] however, includes new profit expected from otherwise impossible ventures into the benefit part of the equation. We have chosen not to make this addition to the ROI formula, because of the difficulty of defining the new profit. This difference should be considered when looking at countermeasures effectiveness.

If the company observes that adding additional countermeasures to the information security infrastructure does not reduce the residual risk to an acceptable level in a cost efficient way, it can choose to invest in cyber-security insurance. However, because of the lack of good actuarial data on which insurance companies can base premiums, they tend to include additional risk factors into their calculations, thus increasing premiums [5].

The company can also change the percentage values of the inputs and see the affects on the outputs and thereby identify areas where they need to spend more money. A 10%

[1]The company must also consider the implications of such a countermeasure on productivity. The second anti-virus may slow down the speed of end-users computers, as more operations have to be conducted due to the secondary anti-virus, and thus negatively affect productivity. Users may also become impatient and attempt to bypass the secondary anti-virus. This would have to be supplemented with additional education and training, thereby increasing costs.

Table 4: **Case study 2: Countermeasure effectiveness**
(All amounts in thousands of dollars.) Note the disproportionate SPVs obtained which indicate that the loss numbers reported by the organization are overly pessimistic. ROIs (not shown) are also disproportionately high.

Countermeasure	Cost	Curr. NPV	NPV w/ r, t
Anti-virus	$1K	$26,154K	$59,702K
Firewall	$0.8K	$18,908K	$43,171K
IDS	$1.5K	$10,808K	$24,676K
Training and education	$3K	$8,112K	$18,520K
Backup server	$2K	-$1,700	-$5,100
Spam filtering	$0.4K	$21,609K	$49,335K

increase in estimated attack frequency, results in an approximately 10% increase in the residual risk. Therefore we can conclude that investing resources into more accurate data collections with respect to attack frequencies would not be cost effective.

3.2 Non-profit organization

This case study is for a charity organization based in the United Kingdom. The currency values have been converted to dollars.

The organization has an annual turnover of $12m and 56 employees. A typical IT employees salary is $60,000, which equates to an IT employee cost per day of $250. The company faces malicious code infections and administrative account compromises.

We compute the residual risk to be $145,578. With 95% confidence, the residual risk should be no more than $232,336, and with 99% confidence our model tells us that over the next year the residual risk will be no more than $261,695.

Our model informs us of the potential effectiveness of additional countermeasures. Using, as in the first case study $r = 0.15$ and $t = 3$ years, and considering countermeasures in isolation, we obtain Table 4. The ROI and NPV for all of the countermeasures is positive, except for the backup server.

We note all numbers in the table are astoundingly high, which indicates the organization is highly sensitive to any change in the security policy. Also, these numbers are due to the high losses as reported by the managers from the organization, compared to the relatively modest turnover. Indeed, according to the values in this table, the mere threat of spam could bring this organization down. This leads us to believe that the self-reported values are overly pessimistic, and illustrates the value of a quantitative analysis of the kind as a "sounding board" when planning a budget. Although the values given are too pessimistic, the respective order of importance of each threat appears to be properly assessed.

The ROI/NPV for the backup server is negative here, because in itself, it does not prevent any attacks; however, it is worth noting that it could be very useful to mitigate (or even completely avoid) information destruction. This again, illustrates the point that the purpose of the tool is to act as a decision support tool, and the decisions are ultimately down to management or the user, who base their decisions on numerous other factors, other than ROI and NPV. These include things such as the profit gained from projects that are made possible because of the countermeasure.

Here again, the attack outcome with the highest expected loss is information destruction. In the case of this organization, administrative account compromise is the sole attack that results in information destruction. Using this information the user can identify potential countermeasures to prevent root compromises. This can include improved IDS and firewall capabilities, activity logging and improved policies with user training.

Table 5: **Case study 2: Expected loss per attack outcome**

Attack Outcomes	EL per Attack Outcome
Information Destruction	$3,437.50
Service (User PC) Unavailable	$962.50
Service (Email) Unavailable	$1,375.00

The user can explore the ROI and NPV for each of these countermeasures by using the same procedure highlighted in the previous case study. We will take a look at improved policies with user training as a possible countermeasure against root attack. Assuming the new countermeasure can reduce the current number of attacks by 50% the number of root compromise attacks would reduce from 10 to 5. This will reduce residual risk to $140,852, that is, a reduction of $4,726. Assuming that instigating the new policy and training users will cost in the region of $4,000, we can see that the countermeasure is cost-effective as the net benefit will be greater than zero. However, it would be better to have a countermeasure that would produce a greater net benefit, and a greater reduction in the residual risk.

Also, in this organization we found that the number of malicious code infection type attacks is quite high (225). This is another area where improvement in preventing losses may be possible. One possibility is to improve the effectiveness of the firewall if a large proportion of malicious code infection attacks that result in root compromises originate from outside the organization is high. Alternatively, the organization could add a secondary firewall. If we explore the idea of making the firewall rules stronger, so that the firewall prevents more of the malicious code infections, we would expect there to be fewer malicious code infections, resulting in fewer root compromises and ultimately a lower residual risk. However, stronger firewall rules would also prevent more legitimate traffic from passing through the firewall, and may impact users negatively. If the user is able to estimate the cost of the strengthened firewall rules, it would be possible to calculate the net benefit and ROI for the countermeasure.

Assuming that the annual negative effects of the stronger firewall are estimated at $20,000, the change in permissions costs $100, and the stronger firewall prevents a further 30% of malicious code infection type attacks, the residual risk will reduce to $77,516. This equates to a benefit of $68,062 ($145,578 - $77,516), a net benefit of $47,962 ($68,062 - $20,100), an ROI of 238% ($47,962 / $20,100), and with cost of capital at 15%, over 2 years, the NPV will be as follows:

$$\text{NPV} = \frac{145578 - 77516 - 20100}{(1+0.15)^2} - 20000 \approx \$16,266 \; .$$

The Value at Risk will now be such that, with 95% confidence the organization will not lose more than $122,365 over the next year from the attacks defined earlier.

Hence, our model tells us that, in the current situation, strengthening the rules to the existing firewall will be a cost-effective loss mitigation strategy. This is however, dependent on the reliability of the inputs to the model.

3.3 Project in multinational insurance company

This case study is for a project within a large multinational insurance company located in Japan. The currency values have been converted to dollars.

The project involves a turnover of $10m and 100 employees. A typical IT employees salary is $60,000, which equates to an IT employee cost per day of $250. The company faces the following attacks: (1) Malicious Code Infections, (2) Account Compromise, (3) Theft, (4) Spam, and (5) Natural Disaster.

Table 6: **Case study 3: Countermeasure effectiveness**

Countermeasure	Cost	ROI	Curr. NPV	NPV w/ (r, t)
Anti-virus	$8K	450698%	$36,057K	$82,298K
Firewall	$10K	503975%	$50,399K	$115,037K
IDS	$10K	472471%	$47,248K	$107,844K
Training and education	$5K	287970%	$14,399K	$32,866K
UPS	$10K	-100%	-$8.5K	-$32K
Server Room – Phys. Sec.	$8K	121681%	$9,735K	$22,209K
Employee Monitoring	$5K	194777%	$9,739K	$22,227K
Active Directory	$10K	21%	$3.6K	-$5K
Backup Server	$15K	68458%	$10,271K	$23,413K
Spam Filtering	$10K	367881%	$36,789K	$83,964K
BCP/DR	$30K	32375%	$9,717K	$22,125K

These result in the following attack outcomes: (a) Information Theft/Disclosure, (b) Information Modification, (c) Information Destruction, (d) Service Unavailable - User PC, (e) Service Unavailable - Email, (f) Service Unavailable - Website, and (g) Legal/Compliance Damage.

Our model predicts a residual risk of $4,521. With 95% confidence, over the next year the residual will be no more than $6,334, and with 99% confidence, the residual risk will be no more than $6,994.

Table 6 shows the ROI and NPV for countermeasures, with $r = .15$ and $t = 3$ years for each of the countermeasures acting alone. The ROI and NPV for all of the countermeasures is positive, except for the UPS. The ROI for the UPS is negative because it does not prevent any attacks. The UPS mitigates the loss from the attack outcome instances. It does not prevent any attacks, therefore it has a negative ROI and NPV.

Based on the Value at Risk figures, the company may decide that the current countermeasures in place are sufficient and any further countermeasures are not needed. This is especially the case as, with 99.9% confidence, given the inputs given, the residual risk for the project will be less than $7,654.

4 Discussion and conclusions

We provide a simple model for quantitative risk analysis of information security, and use this model on three cases studies: a small IT company setting, a non-profit organization, and a project within a multinational insurance company. The model has shown itself to be a useful input to decision-makers in that it has allowed the small IT company to make a decision to introduce a secondary anti-virus and thereby reduce residual risk by $10,000, and calculate that the ROI for the secondary anti-virus would be over 200%.

For the non-profit organization our research has helped the organization to make the decision to strengthen its firewall rules, thus approximately halving the organizations information security related residual risk. It has allowed the insurance company to determine the Value at Risk, giving them a better understanding of the projects risk exposure.

This shows, that used effectively, quantitative models can be an effective decision support tool. User feedback reported that such analysis will now allow them to justify to management countermeasures that they previously wanted to introduce, but for which they were unable to provide a suitable business case.

This work is clearly a long-term research endeavor, of which we have only completed the first step, that is, a methodology definition, and acquisition of initial data. The most interesting part of this research lies ahead of us, in the interpretation of the data. What makes

for instance our third case study to have so little value at risk, compared to our second case study? At first glance, both organizations seem to have some security controls in place, so why do we see such a huge disparity? We are in the process of analyzing this data, and determining if we can derive over-arching security principles tying an organizational structure with possible effectiveness of countermeasures.

Of course, such future work hinges on obtaining even more data to inform an analytic model of countermeasure efficiency. It is our hope that, by making our methodology and its instantiation (Excel spreadsheet) available to the public, we will be able to acquire supplemental data, and foster further discussion between the academic and management communities.

Acknowledgments

We thank our industrial partners for the extensive feedback and data, they provided throughout this research. This work greatly benefited from discussions with Ashish Arora and Rahul Telang at Carnegie Mellon's Heinz School. Mohammed A. Bashir's research was fully supported by a scholarship from the Hyogo Institute of Information Education Foundation.

References

[1] Secure insight analysis, 2007. `http://www.msbai.com/`.

[2] C. Alberts and A. Dorofee. An introduction to the OCTAVE method, January 2001. `http://www.cert.org/octave/methodintro.html`.

[3] A. Arora, D. Hall, A. Pinto, D. Ramsey, and R. Telang. Measuring the risk-based value of IT security solutions. *IEEE IT PRO*, November/December 2004.

[4] B. Blakley. A measure of information security in dollars. In *Proceedings (online) of the First Annual Workshop on Economics and Information Security (WEIS'02)*, Berkeley, CA, May 2002.

[5] L. Gordon, M. Loeb, and T. Sohail. A framework for using insurance for cyber risk management. *Communications of ACM*, pages 81–85, March 2003.

[6] D. Greer, K. Hoo, and A. Jacquith. Information security: Why the future belongs to the quants. *IEEE Security and Privacy*, pages 24–32, July/August 2003.

[7] K. Soo Hoo. *How Much Security Is Enough? A Risk Management Approach to Security*. PhD thesis, June 2000.

[8] J. Meritt. A method for quantitative risk analysis. In *Proceedings of the 22nd National Information Security Systems Conference*, Arlington, VA, October 1999.

[9] T. Peltier. *Information security risk analysis*. CRC Press, Boca Raton, FL, 2nd edition, 2005.

[10] S. Scalet. Risk: A whole new game. *CSO Magazine*, December 2002.

[11] SkyBox Security. Skybox view, 2007. `http://www.skyboxsecurity.com/products/overview.html`.

[12] D. Tan. Quantitative risk analysis step-by-step, 2002. SANS Institute Reading Room paper #849. `http://www.sans.org/reading_room/whitepapers/auditing/849.php`.