Multi-Attribute Risk Assessment

Shawn A. Butler

Computer Science Department
Carnegie Mellon University
Pittsburgh PA 15213

shawnb@cs.cmu.edu

Voice: 412-268-8101
Fax: 412-268-5576

Paul Fischbeck
Associate Professor
Department of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh PA 15217

pf12@andrew.cmu.edu

Voice: 412-268-3240
Fax: 412-268-6938

Abstract

Best practice dictates that security requirements be based on risk assessments; however, simplistic risk assessments that result in lists of risks or sets of scenarios do not provide sufficient information to prioritize requirements when faced with resource constraints (e.g., time, money). Multi-attribute risk assessments provide a convenient framework for systematically developing quantitative risk assessments that the security manager can use to prioritize security requirements. This paper presents a multi-attribute risk assessment process and results from two industry case studies that used the process to identify and prioritize their risks.

1. Introduction

Best practice dictates that security requirements be based on risk assessments; however, simplistic risk assessments that result in lists of risks or sets of scenarios do not provide sufficient information to prioritize requirements when faced with resource constraints (e.g., time, money). In this paper, we show how multi-attribute analysis techniques from the field of decision sciences can be used by security managers to prioritize their organization's threats and in turn their security requirements. The techniques are flexible, systematic, and repeatable and can easily incorporate new threats or changes in the threat environment. With this approach, rationale behind security decisions can be communicated and justified to others in the organization. Uncertainty in the timing and severity of attacks is modeled, and sensitivity of assumptions and subjective estimates are calculated.

1.1 Background

Traditional risk assessments that include asset valuation do not always capture the essence and uncertainty of the underlying risks. For example, it is often difficult for an organization to quantify the damage that a successful attack does to their corporate image. However, this damage to corporate image may be far more important to the organization than the actual loss of revenue caused by the attack or the hours it takes to recover. This evaluation is further complicated when likelihood and severity of the attack (on both revenue and image) are uncertain.

To be effective, security requirements must reflect the risks and priorities of the organization. Only then can information technology managers allocate limited resources appropriately. Multi-attribute analysis techniques provide a convenient framework for developing a quantitative risk assessment that prioritizes both the set of threats and the security requirements.

Multi-attribute risk assessments have many advantages. Most importantly, they allow security managers to identify their organizational risks, express their expectations about the consequences of successful attacks, and provide insights into how the uncertainty of their expectations affect the prioritization of security requirements. In addition, multi-attribute risk-assessments provide a

systematic and repeatable method for evaluating an organization's risks using the best available threat information. As security managers gain better threat information, the risk assessment models can be easily updated with new input data and the marginal effect on security requirements can be measured[3]. The value of a multi-attribute risk assessment is not only in the numbers produced, but also in the insights that security manager's gain during sensitivity analysis and each refinement step of the assessment.

1.2 The Case Studies

This paper presents two case studies that demonstrate how multi-attribute risk assessment techniques can be used to capture a security manager's experience of an organization's risks and then to prioritize the risks and the security requirements. The techniques presented in this paper are based on extensive interviews with security managers and information technology executives at a variety of organizations. The elicitation process was refined so that the necessary threat information could be collected using semi-automated tools, and threat evaluations and sensitivity analyses could be conducted seamlessly.

The first case study risk assessment is from a large commercial organization, which has several information technology centers throughout the world. The second case study risk assessment is from a small hospital information system organization. In all cases, the security managers felt that the method was very helpful in development of a prioritized list of threats that they could use to develop security requirements and justify security budgets. Information technology managers especially appreciated the systematic approach the method offers.

1.3 Roadmap

A brief introduction to multi-attribute analysis techniques and terminology are presented in Section 2. Underlying assumptions and the construction of the multi-attribute risk assessment additive model is shown in Section 3. Section 4 discusses the multi-attribute risk assessment process. Section 5 presents the results of two multi-attribute risk assessments, and Section 6 provides our initial observations about using the process. Finally, Sections 7 and 8 discuss future work and conclusions, respectively.

2. Multi-attribute Analysis

Multi-attribute analysis provides a convenient framework for developing a quantitative risk assessment that results in a set of prioritized risks. Though there are different forms of multi-attribute models (e.g., linear, multiplicative, power), the work presented here relies on an *additive value model.* This model provides an intuitive mechanism for ranking each threat and allows the security manager to express outcomes that involve a wide variety of concerns (e.g., lost income, reputation, market share) in unified, non-economic terms.

2.1 Why Multi-attribute Analysis

Multi-attribute analysis techniques help decision makers evaluate alternatives when conflicting objectives must be considered and balanced and when outcomes are uncertain[2]. Though not demonstrated here, this approach can also integrate a security manager's risk attitudes (i.e., risk aversion or risk seeking). Once constructed, this multi-attribute analysis framework also provides the basis from which the security manager can systematically evaluate alternative risk-mitigation strategies.

The risk-assessment process can be structured as a *multi-objective*, *compensatory* decision problem. *Multi-objective* decisions are those whose consequences can be characterized using multiple attributes. As part of the risk-assessment process, security managers identify their

organization's threats and the potential consequences or outcomes from successful attacks. Threats are defined as events, such as denial of service attacks, procedural violations, IP spoofing, etc., which could lead to an information system compromise. An attack (a) is an instance of a threat that results in an information system compromise that has an outcome (O_a) of one or more consequences (X_i). For example, a system compromise may ultimately result in lost revenue (X_1), public embarrassment (X_2), lost productivity (X_3), and damaged corporate image (X_4).

The types of consequences (X_i) that can result from an attack constitute outcome *attributes* in the multi-attribute risk assessment and the actual attribute damage from an attack is the attribute's value (x_i). Therefore, each attack outcome can be described as a vector of attribute values $O_a(x_1,x_2,x_3,x_4)$. Similar attacks can have many different outcomes because each attack instance results in different consequences. For example, a denial of service attack could result in a few hours of lost productivity to many hours of lost productivity depending on the nature of the attack.

Compensatory decisions permit tradeoffs among attributes. The multi-attribute risk assessment is characterized as compensatory because managers are often willing to trade one outcome for another. For example, some organizations are willing to trade some amount of revenue to avoid public embarrassment. Security managers in these organizations would prefer to invest significantly in security technologies that reduce the risk of attacks that are likely to result in public embarrassment. Other organizations may be much less willing to do so.

2.2 The Additive Value Model

Framing a risk assessment as a multi-objective, compensatory decision problem allows one to use an additive value model[4][6]. The additive value model offers a simple way of evaluating multi-attribute alternatives. The additive value model relies on the additive value function. The general form of an additive value function is:

$$v(x_1, x_2, \ldots, ,x_n) = \Sigma_{i=1,n} w_i v_i(x_i)$$

where $v_i(x_i)$ is a single-attribute value function defined over levels of x_i, and w_i is a scaling constant that weights the value function for attribute value x_i. Constructing an additive multi-attribute value function involves five steps:

- Check additivity assumptions to see if the additive form is valid
- Assess the single-attribute value functions $v_1, v_2, \ldots, v_n$
- Assess the weighting factors $w_1, w_2, \ldots, w_n$
- Compute the value of each alternative and rank alternatives
- Conduct sensitivity analysis to see how sensitive the ranking is to model assumptions

The next two sections describe in more detail the additive multi-attribute value model and risk assessment; however, intuitively, the single-attribute value function ensures that outcome attribute values can be summed together using the weights that reflect the assessed preferences (e.g., security manager).

3. Multi-attribute Analysis and Risk Assessments

Assume that a security manager has identified four outcome attributes of concern given a successful attack: 1) lost productivity, 2) lost revenue, 3) damaged public reputation, and 4) additional regulatory penalties. Additional regulatory penalties are the increased administrative

burdens that an organization suffers from an external oversight agency because of a security compromise. These were assessed using an open protocol that gave the security manger a broad set of attributes that have been shown to be important over many interviews and organizations. Additional concerns can be added as needed.

3.1 Checking Additivity Assumptions

Before the additive value model can be used in the risk assessment, certain relationships must hold. The additive value model is valid if *transitivity*, *preferential independence*, *tradeoff independence*, and *difference independence* conditions exist among the attributes. Although it is not possible to prove the requisite additivity assumptions hold in every case, there is strong evidence that even when there is not complete independence, the additive value model provides close approximations to "pure" additive value functions.[9]

3.1.1 Transitivity

The transitivity condition holds if O_1 is preferred to O_2, and if O_2 is preferred to O_3, then O_1 is preferred to O_3 (where O_1, O_2, and O_3 are outcome vectors). For example, if an organization has only one outcome attribute, lost revenue, and each outcome represented increasing amounts of lost revenue. Then it is reasonable to assume that the security manager would prefer less lost revenue than prefer more lost revenue in every case. Transitivity is a normative assumption that any rational decision maker follows.

3.1.2 Preferential Independence

Preferential Independence exists if the decision maker's preference ranking for one attribute does not depend on fixed values of other attributes. If $X_1,\ldots,X_n$ represent decision attributes, then attribute X_1 is said to be independent of the remaining attributes if the decision maker's preference ranking of different levels of X_1, holding $X_2,\ldots,X_n$ constant, does not depend on where we hold $X_2, \ldots, X_n$ constant.. For example, if X_1 is lost productivity, X_2 is lost revenue, and X_3 is damaged public reputation, then less lost productivity is still preferred to more lost productivity for any fixed combination of lost revenue and lives lost. Although this is not a normative assumption, there does not appear to be any reason to believe that it does not hold for computer security problems. Counter examples do exist in other domains (e.g., city preference is likely to depend on ones salary).

3.1.3 Difference Independence

Difference independence goes a step further than preferential independence. The additive model assumes that a decision maker can rank order the differences in value within an attribute. Difference independence requires that the ranking of the differences in values within the attribute do not change given fixed levels of outcomes in other attributes. For example, in the risk-assessment levels of public reputation and additional regulatory penalties are represented using a 7-point Likert-type scale[1], then the scale intervals should hold no matter the levels of the other attributes.

3.1.4 Tradeoff Independence

Tradeoff independence requires that tradeoffs between two attributes, holding all other attributes fixed, do not depend on where we hold the other attributes fixed. If $X_1, X_2, \ldots, X_n$ are a set of three or more attributes, then tradeoffs for every pair of attributes, X_1 and X_2, do not depend on where we hold X_3 fixed. Although tradeoff independence is also not a normative assumption,

---

[1] A Likert-type scale is an interval scale where the intervals between statements are meaningful. For the outcome attributes Damaged Public Reputation and Additional Regulatory Penalties examples in this paper, a 7 point Likert-type scale was used with 1 meaning *no* or *negligible* damage and 7 meaning *severe* damage.

there is no reason to believe that tradeoffs between two outcomes from successful attacks are affected by the values of other outcomes.

3.2 Assess the Single Attribute Function

The second step in constructing an additive multi-attribute value function is to assess a single-attribute value function for each attribute. The purpose of this function is to reflect preferences for outcomes over the relevant range for each attribute. We standardize the results of a single-attribute value function to a 0-1 scale to eliminate computational problems caused by the different units of measure. For the risk assessments described in this paper, we used an increasing linear function to reflect the consequences and to normalize the attributes. The form of the function is:

$$v_j(x_{ij}) = x_{ij}/x_j^*$$

where x_{ij} is the i^{th} attribute value of the j^{th} attribute and x_j^* is the maximum value for that attribute. This ensures that $0 <= v_j(x_{ij}) <= 1$, and as $v_j(x_{ij})$ approaches 1, the consequence is more severe. Though we successfully used this simple linear function initially for simplicity, future research might indicate where a convex or concave function would better describe the security manager's preferences. If the interview process reveals a different relationship, then an analyst should use other forms of monotonically changing functions (e.g., convex or concave).

3.3 Assess Weighting Factors

The third step in constructing an additive multi-attribute value function is to assess the attribute weighting factors. These weights permit trade-offs to be made between the attributes. Although several weighting elicitation techniques have been developed [8][1][5][7], the Swing-Weight Method is demonstrated here. It is easy to use and the security managers found it cognitively appealing.

In the *Swing-Weight Method,* the analyst asks the decision maker to consider a hypothetical situation, where the security manager discovers a new type of threat. This threat results in the worst level of damage for each attribute. The analyst gives the decision maker the option of improving the hypothetical outcome by changing one attribute to its best level. The first chosen attribute is considered to be the most important and the one that matters most to the decision maker. In succession, the decision maker improves each attribute until all attributes are ranked. Next, the decision maker assigns a value of 100 to the most important attribute and values the remaining attributes in relative importance to the first attribute. The actual weights are determined by dividing each of these values by the sum of all the values. The resulting weights sum to 1.0. Table 3-1 shows the results of using the swing-weight method for the example.

Table 3-1 Outcome Attributes

Outcome Attribute	Rank	Assessed Preference	Weight
Lost Productivity	1	100	.42
Public Reputation	2	80	.33
Regulatory Penalties	3	40	.17
Lost Revenue	4	20	.08

3.4 Compute Value and Rank Alternatives

The fourth step in constructing the additive multi-attribute value function is to compute the relative ranking of the alternatives. For the risk assessment, this means computing the *Threat Index*, which is used to rank the threats. The threat index (TI) captures the relative importance of each type of attack. For the risk assessment, decision makers are assumed to be risk neutral and that utility functions do not have to be assessed. Techniques for integrating preferences of risk-averse decision makers have been developed[6], but can be

complex to implement. The threat index (TI) for each type of attack (*a*) is computed using the following equation:

$$TI_a = Freq_a * (\Sigma_{j=attributes} w_j * v_j(x_{aj}))$$

where w_j is the attribute weight and x_{aj} is the "most likely" outcome attribute value for the attack.

Table 3-2 shows the data and threat index of three threats (Procedural Violations, Theft, and Virus) as an example. The weights (*w*) are shown for each outcome attribute and the second column shows how often the security manager expects an attack to occur. The left-hand column under the outcome attributes shows the most likely consequence of an attack. For example, the security manager expects a procedural violation to occur 4,380 times per year and with each instance, the attack is likely to result in approximately $2 of lost revenue, have a *slight* ('2' on the Likert Scale) impact on the organization's reputation, and lose 2 hours of productivity. The right-hand column under the outcome attributes shows the normalized values from the attribute value functions. The values in the TI column are dimensionless units, but the TI indicates the relative significance of each threat.

Table 3-2 Outcome Attribute Values and Threat Frequencies

		Outcome Attributes								
Threats	**freq/yr**	**Lost Revenue**		**Reputation**		**Lost Productivity**		**Reg. Penalties**		**TI**
		***w*=.08**		***w*=.33**		***w*=.42**		***w*=.17**		
Procedural Violation	4,380	$2	.0002	1	.25	2hrs	.0083	0	0	**376.69**
Theft	24	$182	.0152	2	.5	1hrs	.0042	2	.67	**6.75**
Virus	912	$0	0	0	0	3hrs	.0125	0	0	**80.03**

Once the TI's are computed, the security specialist should be shown the results so that counter-intuitive results can be investigated. Adjustments to the assessed values are permissible if the results uncover errors or provide the decision maker with a clearer idea of the process. Tweaking the input values to obtain preconceived results should be avoided.

3.5 Conduct Sensitivity Analysis

The final step in constructing the additive multi-attribute value function is to conduct sensitivity analysis. The purpose of the sensitivity analysis is to determine how sensitive the analysis is to the security manager's range of uncertainty about key variables. We examine three components in the risk assessment analysis: 1) attribute weights, 2) estimated frequency of attacks, and 3) the attribute values. During the elicitation process, the security manager provides an upper and lower bound to his estimates, in addition to the "most likely" values. These bounds are used to construct probability density functions that are used to create input to simulations. Although these simulations and sensitivity analysis are described in further detail in the next section, the primary benefit of sensitivity analysis is to show how the security manager's uncertainty affects each threat's prioritization and quickly find inconsistencies in the results. In addition, the simulations allow the security manager to conduct "what-if" analysis to test different threat assumptions.

4. Case Study Process

Each risk assessment consists of four sequential steps. Although the participants of a risk assessment vary among organizations, the multi-attribute analyst/interviewer and the organization's lead security manager (or specialist) are the key participants. The analyst is responsible for facilitating the interviews, recording the information provided by the security manager, analyzing the results and conducting sensitivity analysis. The security manager usually relies on other security specialists or information managers within the organization to provide specific information about the threats. This section describes the case study process, which relies on the additive model described in the previous section.

4.1 Initial Threat Definition

During the first step, the security manager determines which threats are potential risks to the organization and orders the threats from greatest to least concern. Recall that we defined *threat* as an event that could lead to an information system security compromise. In contrast, a *vulnerability* is a defect or flaw in the system; therefore, threats exploit vulnerabilities. The security manager can usually fix or patch vulnerabilities once discovered, but most organizations develop security requirements to address situations in which someone takes advantage of unknown vulnerabilities. i.e. a threat.

The analyst provides an initial set of threats to help initiate the risk assessment process, but the security manager refines and orders the list. Although we derived the initial list of threats from actual risk assessments collected from organizations, some threats are not appropriate for all organizations so the security manager tailors the list to meet his organization's risks. Of course, the security manager can use any of his organization's existing risk assessments to help derive the appropriate list of threats. The result of the first step is an ordered list of threats that represent the organization's security concerns.

4.2 Outcome Attributes

In the second step, the security manager provides the risk assessment *attributes*. Recall from Section 3 that a risk assessment attribute is a potential consequence of an attack, such as damage to corporate image, lost revenue, or lost productivity. Once the security manager establishes the top 3 or 4 risk assessment attributes, then (s)he estimates the distribution of attack frequencies and outcome attribute values for each threat. This assessment can be simple with estimates of the most likely value, and upper and lower bounds.

4.3 Threat Indexes

In the third step, the security manager ranks each attribute and an initial *threat index* is determined for each threat using an automated tool to construct and conduct attack simulations. Using the security manager's estimated attack frequencies and expected outcomes, we construct a probability distribution that uses the expected values as the distribution mean, and the low and high outcome values as lower and upper bounds on the distribution. For example, using the frequency data for the Procedural Violation threat from Table 3-2, the analyst creates a Normal Distribution using 4,380 as the mean for distribution, a standard deviation of 440, and 0 and +infinity as the upper and lower bounds of the distribution. Figure 4-1 shows the Frequency distribution curve for the Procedural Violations threat.

Figure 4-1 Procedural Violations Probability Distribution

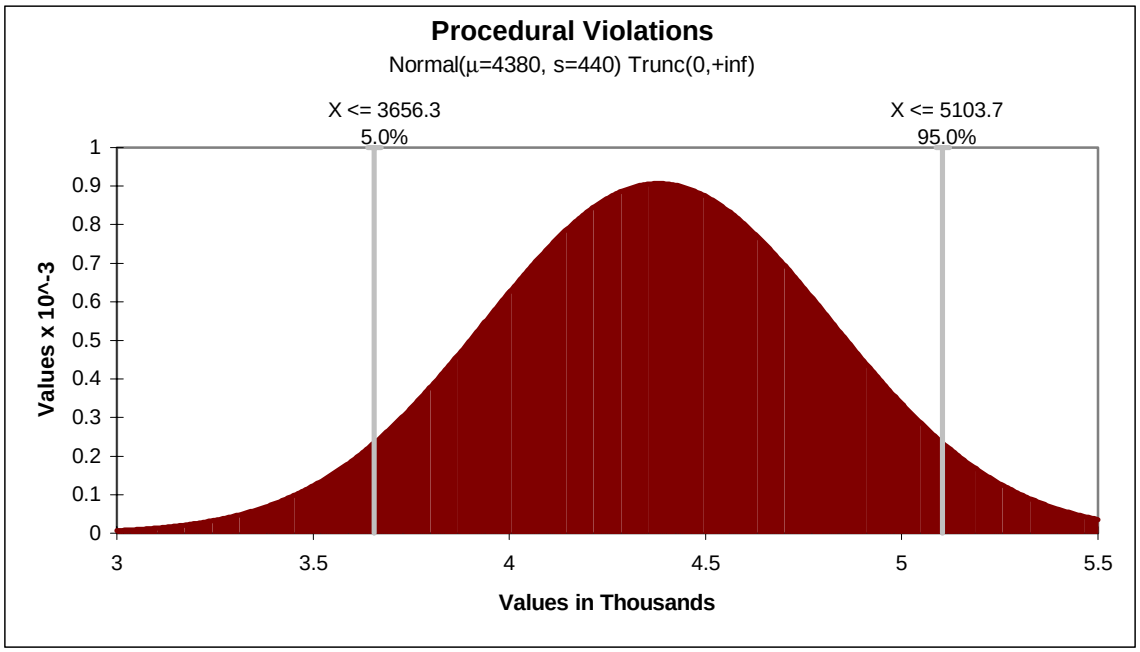

Although the security manager provides the expected frequency or most-likely outcome, we initially set the standard deviation as half of the difference between the lower bound and the expected value. In Figure 4-1, the lower bound is 3,500 so the standard deviation is 440. In addition, the probability distribution is truncated at zero, since that is the lowest possible value for an attack frequency or outcome, but not truncated at the upper bound. The analyst can adjust these values if necessary, and security manager can conduct "what-if" analysis to see how changes in the distribution parameters would affect the threat outcome.

After defining a probability distribution for each threat, we conduct at least 1,000 simulations. A simulation consists of selecting frequency and attribute values for each threat consistent with their distribution curves and calculating the threat index. After running all simulations, the tool computes the average simulation threat index for each threat and ranks the threats.

4.4 Sensitivity Analysis and Refinement

In the final step, we compare the results of the multi-attribute risk assessment with the initial ordering. We conduct sensitivity analysis and present the results to the security manager so that significant differences can be understood and resolved. Significant differences are those risks that differ by greater than five ranks between the security manager's initial ordering and the multi-attribute risk assessment results. For example, if the security manager placed a threat 10^{th} in his initial ordering, but the multi-attribute risk assessment ranked it 16^{th}, then the security manager and the analyst would review the initial ordering and the input data to understand the difference. Of course, any differences in order can be explored.

The most valuable aspect to resolving the differences is the insight it gives the security manager. In order to resolve ordering differences, the security manager can rethink his initial ordering, threat frequencies and outcome values, or the attribute rankings. The security manager can also conduct "what if" analysis to see how different input data affects the threat ranking. This helps the manager understand how the variability around his uncertainty might affect the prioritization of threats. The security manager gains confidence in the results because he can see how his estimates produced the threat priorities and he can determine how sensitive the results are to the

uncertainty of his estimates. In addition, the security manager can focus his/her attention on those aspects of the risk assessment that matter most in the prioritization of the organization's threats.

5. Security Requirements

Once the risk assessment is completed, the security manager can determine which risk mitigation strategies are appropriate for each threat. The security manager can rank each of the strategies according to how effective the security manager believes it is against its risk assessment. For example, assume that a security manager identified antivirus software, mobile code scanners, and secure email as risk mitigating for viruses. Furthermore, if the risk assessment determined that viruses are one of the most significant threats to the organization, then one or more of these technologies should be a high priority security requirement. In contrast, if the risk assessment determined that denial of service attacks were not a significant threat to the organization, then denial of service security requirements would also rank low.

One of the greatest advantages to the multi-attribute risk assessment, and subsequent prioritizations of security requirements, is that information system managers can understand and trace the justification for requirements and the basis of security resource allocation. When the underlying assumptions to the risk assessment change, i.e. the frequency and outcome estimates, information system managers can see the impact of those changes on the security architecture.

6. Case Study Results

The objective of the multi-attribute risk assessment is to provide the security manager with insight into the significance of the threats so that security requirements can be developed to address the organization's risks. This section discusses how well the multi-attribute analysis correlates with the security manager's initial threat ordering and provides some feedback about the case study security managers' insights that arose from the risk assessment process. Although the results presented in this section are from two multi-attribute risk assessments, similar results have been achieved in other case studies.

6.1 The Threats

The commercial organization's security manager selected and prioritized 27 threats in the first step of the multi-attribute risk assessment. The hospital security manager selected and prioritized 15 threats. Both organizations' security managers had staff security specialists participate in the selection and prioritization process, but neither organization had ever completed a formal risk assessment and did not maintain historical security incident records. Even if an organization maintains security incident records, the data is usually not sufficient to address all the threats necessary for the risk assessment, especially outcome data. Rare attacks that have not historically happened must be considered.

The commercial organization identified Damage to Public Image, Damage to Customer Relationships, and Lost Revenue as the three most important concerns from a security compromise. In actuality, most threats resulted in lost productivity rather than lost revenue, however the security staff converted lost productivity into lost revenue by estimating the amount of lost hours and multiplying the hours by an average employee hourly rate. Although this is a standard technique in many other types of risk assessment, it failed to capture the essence of the outcome.

In this particular organization, a threat that resulted in lost revenue was more significant than a threat that resulted in lost productivity, even if they were equivalent in monetary terms. In general, the multi-attribute risk assessment ranked threats that resulted in lost productivity higher, i.e. more significant, than the security manager's initial subjective ordering. While differences between lost productivity and lost revenue may not apply in all organizations, the results in this case study show one of the problems of reducing all outcomes to economic terms.

Although the threat index is not a precise measure of risk, it does provide a relative indication of importance to the security manager. In both case studies, the multi-attribute risk assessment method determined that only a few threats constitute the most significant threats to the information system. Figure 5-1 shows the hospital risk assessment results. Three threats, Alteration[2], Virus, and Compromising Emanations[3] constitute 70% of the total threat index. In other words, these three threats are the greatest concern to the hospital and the other twelve threats are significantly less important.

In the commercial case study, one threat, virus, dominated all others. One possible explanation for this was that the security manager had spent considerable effort in the past six months addressing a virus problem and had accumulated a lot of information about the frequency of viruses in their organization. The model ranked it number one, consistent with the security manager, but the threat index should have been much lower relative to the other threats.

Figure5-1 Threat Index Percentages

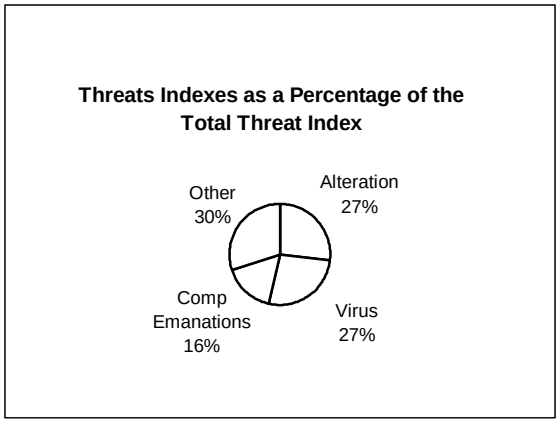

6.2 Ordering Comparisons

As previously mentioned, in the first step of the multi-attribute risk assessment process, the security manager provides an initial ranking of the threats. To determine model agreement, we compared this initial ranking with the results of the multi-attribute risk assessment by calculating the rank correlation. The correlation coefficient for the commercial security manager's ordering and the multi-attribute risk assessment was .19. The correlation coefficient for hospital's threat ordering and the risk assessment was .53.

After reflection and model refinement, these correlations increased in both case studies. In the commercial case, the security manager changed several of his initial threat rankings based on the results of the multi-attribute risk assessment. In addition, the multi-attribute risk assessment highlighted the internal threats to the organization that justified countermeasure strategies that reduced internal risks. In the hospital case, the security manager felt the initial threat ranking was accurate and their assessment of threat frequency and outcome were wrong. In addition, further evaluation of their top threats showed that the organization needed more risk mitigation strategies that could help detect internal security compromises. The final correlation coefficient for the commercial case study was .86 and the final correlation coefficient for the hospital orderings was .81.

One limitation of comparing the security manager's threat prioritization with the risk assessment results is that we compare two different types of scales. Although the risk assessment produces threat indexes, which are relative rankings of threats the security manager provides only an ordering. The security manager's ordering provides little insight into the magnitude of the differences. For example, the security manager can rank two threats 8th and 9th, but the difference

[2] Alteration is the intentional modification, insertion, or deletion of data or lines of code, the compromises the auditability, confidentiality, recoverability, availability, or integrity of data or application.

[3] Compromising emanations are the unintentional data-related or intelligence-bearing signals that, if intercepted and analyzed, could disclose classified or sensitive information being transmitted or processed.

between the two could be minimal or quite large. One way to resolve this problem is to have the security manager use a ratio scale initially so that the differences between ordering has meaning

7. Observations

Both organization security managers thought that the initial list of threats we provide helped them think about risks that they had overlooked. Clearly security managers were more familiar with some threat than others, but the multi-attribute risk assessment process required them to think about why a particular threat ranked high, rather than reacting to the latest security trade journal article. In the end, the security managers were very satisfied with the results.

In addition to unfamiliarity with threats, security managers usually don't have good data to support all threat frequency estimates or outcomes, especially on attacks they haven't yet experienced. If the security manager can't find industry data to support the estimates, they can estimate a wider range of upper and lower boundary estimates or conduct "what-if" analysis to see how the different estimates affect the threat prioritizations. If security managers obtain better data later, they can quickly see how the new data impacts the prioritizations.

We conducted the two case study risk assessments using interviews, but the process can be time consuming for security managers who may not be able to dedicate several hours for each interview. Therefore, we developed a questionnaire that allowed the security manager more flexibility in providing the frequency and outcome data. One advantage of the questionnaire is that the security manager can easily compare threat frequencies and outcomes to ensure consistency among similar threats. Overall, we prefer the interview process over the questionnaire because there is a lot of information brought out during the interview that helps during the refinement step.

8. Future Work

Undoubtedly, the most often asked question from the security managers is "How do we compare to other organizations?" For example, the hospital's security manager would like to know what other hospital security manager's perceive as their most significant threats based on their threat frequencies and outcomes. As we conduct more multi-attribute risk assessments, we will establish a database so that organizations can get insight into similar organizations' risks.

Another interesting research question is how the multi-attribute risk assessment results will change as organizations update their risk assessments. We are confident that since the process is structured, security managers will be able to fine-tune their assessments because they will start to collect information about attack outcomes. In addition, we could compare security requirements and expenditures against threat prioritizations to see if limited resources are allocated against the highest priority requirements.

9. Conclusion

The feedback from security managers has been very positive, and they quickly understand the value of the sensitivity analysis, which give the results much more credibility with their information system managers. As we conduct more risk assessments, we will develop a database of threat information that will help security managers calibrate their own estimates and fill in where security specialists lack experience. Finally, security managers have a systematic and repeatable method they use to develop and prioritize requirements, which can consistently evaluate changes in the threat environment.

10. References

[1] Borcherding, K., T. Eppel, et al. (1991). "Comparison of Weighting Judgments in Multiattribute Utility Measurement." Management Science **37**(12): 1603-1619.

[2] Bunn, Derek W. "Applied Decision Analysis", McGraw Hill, 1984.

[3] Butler, Shawn A. “Security Attribute Evaluation Method: A Cost Benefit Approach.” 24th International Conference on Software Engineering Proceedings, May 2002: 22-240.

[4] Edwards, W. (1977). “How to Use Multiattribute Utility Measurement for Social Decision-Making.” IEEE Transactions on Systems, Man, adn Cybernetics **7**(5): 326-340.

[5] Fischer, G. W. (1995). “Range Sensitivity of Attribute Weights in Multiattribute Value Models.” Organizational Behavior and Human Performance **62**(3): 252-266.

[6] Keeney, R.L. and Raiffa, H. *Decisions with Multiple Objectives*. New York:Wiley. 1999.

[7] Stillwell, W. G., D. A. Seaver, et al. (1981). “A Comparison of Weight Approximation Techniques in Multiattribute Utility Decision-Making.” Organizational Behavior and Human Performance **28**: 62-78.

[8] von Winterfeldt, D. and W. Edwards (1986). Decision Analysis and Behavioral Research. New York, NY, Cambridge University Press.

[9] Yoon, K. Paul and Hwang, Ching-Lai *Multiple Attribute Decision Making: An Introduction*, Sage Publications, 1995.

11. Acknowledgements

This research was supported in part by the National Science Foundation under Grant CCR-0086003 and by the Software Engineering Institute, under the Networked Systems Survivability Program