

Decision Analysis by Proxy for the Rational Terrorist

Heather Rosoff¹ and Richard John^{1,2}

¹Center for Risk and Economic Analysis of Terrorism Events

²Department of Psychology

University of Southern California

Rosoff@usc.edu; Richardj@usc.edu

Abstract

This paper describes a methodology for representing terrorist preferences for alternative modes of attack on the US. The model includes multiple and conflicting objectives related to the attack, attitudes toward risk, trade-offs among various objectives, and uncertainty about the success of particular attack strategies. The methodology utilizes judgments from experts who are knowledgeable about terrorist motivations and beliefs to provide assessments of relevant terrorist leader uncertainty, as well as uncertainty in their own knowledge about the terrorist beliefs and values. A multi-attribute utility model is embedded within a simulation model that generates risk profiles for each attack mode strategy, as well as estimated probabilities that a particular terrorist leader will select each attack strategy over a fixed time horizon. Results of a demonstration with 4 proxies and 9 different attack mode strategies are presented.

1 Introduction

Understanding the objectives and motivations that drive terrorist group behavior is critical. Current methods for terrorism risk assessment focus on target vulnerability, terrorist capability and resources, and attack consequence. What many researchers have yet to consider is the influence of terrorist group values and beliefs in deciphering the root cause of their militant behavior. This understanding has the potential to contribute to probabilistic estimates of terrorist threats.

Using a value-focused decision framework [Keeney, 1992] we refer to as “proxy utility modeling” we assess how the values and beliefs of terrorist leaders might influence the selection of an attack strategy. We then use a random utility modeling approach to compare the risk profiles of alternative attack strategies and estimate the relative likelihood of a terrorist leader selecting a particular attack strategy.

Since we cannot collect information directly from terrorists for the model development, individuals who have studied the general topic of terrorism as well as Islamic terrorist groups (such as Al Qaeda) were asked to act as Al Qaeda terrorist leader value experts (proxies). These proxies included people holding positions as former intelligence specialists, policy analysts and researchers

familiar with the study of terrorism and associated events, and former residents of Middle Eastern Islamic countries familiar with the perspective and motivations of Islamic terrorist organizations. We expected some proxies to be more informed, others to disagree, and of course none are perfect. The diversity across proxies was critical for the assessment of various perspectives on how terrorist leader values and beliefs motivate the attack strategy decision process.

The next section of this paper describes the use of risk assessment in the evaluation of the terrorist threat. Section three summarizes an analysis of terrorist attack alternatives and describes the MAU assessment and elicitation procedures employed. Section four presents results from the MAU model and the implications of the findings upon the overall probability of attack.

1.1 Terrorism Context

Several applied case studies exist assessing the risk of disruption to the electrical grid [Simonoff *et al.*, 2007], the effects on the U.S. economy of a seven-day shutdown of the commercial aviation system following an attack [Gordon *et al.*, 2005], and a dirty bomb attack upon the ports of Los Angeles and Long Beach [Rosoff and Winterfeldt, 2007], to name a few. These reports offer tremendous insight into the technical and resource (manpower) capabilities of the terrorist organization, the relative feasibility of carrying out and defending against the said attack, and the economic, health and psychological consequences that might ensue. To fully assess the threat of terrorism, studying potential terrorist attack targets alone may not be the most effective counter-terrorism strategy. This paper considers incorporating the probability of an attack being selected for execution into the analysis.

A terrorist organization’s commitment to an attack’s execution is part of a complex decision process. Much like the Department of Homeland Security makes decisions on national counter-terrorism policies [Keeney, 2007], terrorists must decide upon the best attack strategy given their perceived security needs. If certain beliefs or motivations weigh heavily on a leader’s decision making process, then certain attack types may have an increased likelihood of occurring. This chapter describes an approach to modeling the decision problem of a terrorist leader. By modeling the terrorist leader mindset, additional information

is acquired about the decision making relative to attack selection.

In the study of terrorism, estimating the likelihood of alternative outcomes is complicated by the nature of the attack type being unpredictable in terms of the time and location of the event. Other various disaster situations, whether they are technologically, manmade or naturally based have faced similar predictive challenges. Yet, researchers still attempt to characterize the probability of these events. They conduct geological studies to evaluate earthquakes, oceanographic studies to understand hurricanes and risk studies to assess the threat of industrial accidents.

The study of terrorism is further complicated by the fact that it is difficult to identify terrorist leader preferences. Collecting data on terrorist leaders' values and beliefs is a formidable task, given the sensitive nature of the information and limited number of public resources. For example, a suicide truck bombing might be the most feasible attack alternative, but there is uncertainty about whether the attack's outcome meets the objectives of the terrorist leader's values and beliefs. Alternatively, a dirty bomb attack might be desirable to a terrorist, but this is conditional upon the success of acquiring the radioactive material. Any decision model build around terrorism will have to account for the uncertainty that the terrorist has about the alternatives, as well as the uncertainty of the analyst's assumption relative to the terrorist's preferences.

Terrorist attacks are created and caused by human agents and thus, extremely dynamic in nature. Knowledge about the functionality of terrorist leaders, their organizations, and their capabilities, is perpetually evolving and difficult to acquire. There is a need for a systemic approach to assessing the uncertainty associated with the decision making threat posed by terrorist leaders. A greater understanding of the opponent's objective function may give some direction as to the probabilities associated with different attack types. This article describes the construction of a value model used to decompose the decision of a terrorist leader. To accomplish this, we seek to utilize and expand upon the multiple objective decision analysis approach [Keeney, 2007], [Keeney and Raiffa, 1976].

1.2 Proxy Multiple Objectives Value Modeling

The primary objective of the MAU model is to use proxies' interpretation of Al Qaeda leaders' beliefs and motivations to determine what might be their preferred attack type. The construction of the model involves six primary steps:

- (1) Select the fundamental objectives for values analysis,
- (2) Identify and define attributes for the fundamental objectives,
- (3) Assess the risk preferences for the attributes,
- (4) Define the value tradeoffs that prioritize the different objectives and attributes,

- (5) Specify the uncertainties of the attributes and all model parameters,

- (6) Use Monte Carlo simulation to obtain risk profiles (CDFs) for each alternative, and probabilities that the utility of each alternative attack strategy is its maximum.

As previously noted since we cannot directly collect information from terrorists, proxies were used for elicitation. As a result, probability distributions were assessed over attribute scale scores, utility function parameters, and trade-off (weight) parameters to address any uncertainty in terrorist prediction of future outcomes and about the proxy's uncertainty about terrorist beliefs and uncertainty about the event tree probabilities. In addition to capturing uncertainty about proxy terrorists' values and beliefs, there is also considerable uncertainty associated with the feasibility of attack execution. Analysts incorporated an event tree into the MAU model to account for variations in attack execution, and to decipher how, if at all, variability in success probability might impact the proxy's preference for an attack alternative. Overall, this value model design characterizes what we think Al Qaeda leaders believe and ultimately, will provide us with insight into making better decisions about defending against the terrorist threat.

2 Methodology

This section describes the development of the value model of proxy Al Qaeda experts' preferences to help represent the values and beliefs behind terrorist leader attack strategy selections. An introduction to the proxy decision maker is initially provided, then a set of objectives and attributes are defined, and lastly, the assessment of a utility function over these attributes is developed.

2.1 Decision Maker and Context

To identify objectives critical to a decision analysis, it is important to assess information about the decision makers' value and beliefs; in the context of this study, the decision maker is a hypothetical leader of Al Qaeda. While the actions carried out by terrorist organizations are interpreted by many to indicate that terrorist leaders are irrational and their decision processes are devoid of rationale, published writings suggest otherwise [Sprinak, 2000], [Victoroff, 2005]. Dating back to the early 19th century, a rational justification for terrorism was made in the context that violence was recognized as a means to an end [Crenshaw, 1995]. Terrorists pursued goals recognizing that the consequences might be grim, yet they had a practical determination. Under such pretenses, terrorists were assumed to operate as a collective unit that required a high level of organization and careful planning to succeed [Rosoff and Winterfeldt, 2007].

The idea of a terrorist group being rationale also translates into how terrorist leaders make logical and

strategic decisions. In this framework, terrorism is perceived as an instrumental activity designed to achieve a set of goals. Like any such decision, the terrorist leader evaluates a decision by looking ahead and evaluating consequences, which in this case refers to the decision to commit a terrorist act and the nature of the attack selected. Much like any other major business or social development decision, a terrorist leader attempts to maximize expected returns while minimizing the expected costs in terms of lives and dollars spent.

2.2 Objectives Hierarchy

The basic way to derive objectives is to start by asking individual decision makers about the meaning and reasoning behind what drives a terrorist organization to commit acts of terror. Then for each fundamental objective, discussions about mechanisms for obtaining them are ensued. To elicit this information, the proxy terrorists were first interviewed individually, and then from the individual assessments a union of the provided objectives was developed. The findings were organized into the objectives hierarchy illustrated in Figure 1.

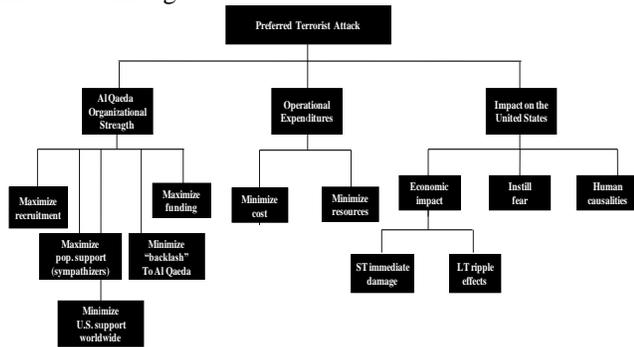


Figure 1. Terrorist Leader Objectives Hierarchy

It was determined that an Al Qaeda leader’s perceived primary objectives fall into three categories: (1) Maintaining Al Qaeda’s organizational strength, (2) managing Al Qaeda’s operational expenditures, and (3) ensuring Al Qaeda has an impact upon the U.S. Further investigation into the primary objectives resulted in a compilation of attributes, or sub-objectives, that are used to evaluate and measure the aforementioned primary objectives. For example, one objective of a terrorist leader is to continue the formation of training camps. This is a means to the larger objective of maintaining internal organizational strength in the event of an attack. This is also a means to contributing to the fear of terrorism in the U.S., as the suggested existence of training camps implies the Al Qaeda threat is in fact a reality. Overall, the attributes identified were of a health, economic or psychological nature.

2.3 Attack Alternatives

Proxies were asked to suggest attack modes (strategies) that Al Qaeda leaders would contemplate to achieve ideal attack feasibility. These suggestions were collected during open

dialogue about the terrorist organization’s general objectives and operations. Analysts opted for this style of interview format to ensure that each proxy formulated his suggestions from whatever considerations, such as attack feasibility, attractiveness, consequences and so forth, he so chose.

Table 1 is the compiled list strategies. The strategy list resulted in a range of attack modes including explosive, nuclear and biological alternatives. The “no attack” alternative was included as an obvious status quo alternative for comparison; some proxies also indicated doubt that Al Qaeda leaders were actively preparing to execute attacks at this time. They posited that the organization might be directing their resources to internal development or toward the execution of attacks outside of the U.S.

Attack Alternatives
No attack (baseline)
IED* in the engine room of naval vessel
Explosion resulting in dam failure
MANPADS** attack on an airplane
Portable nuclear bomb in a major city
Explosion in mass transport(s)
Release of anthrax (movie or sports arena)
Detonation of a dirty bomb
Smallpox release in a major city

*Improvised explosive device
 **Man Portable Air Defense Systems

Table 1. Attack Alternatives

The attack strategies were deliberately broadly defined as strategic initiatives that allowed each proxy to flexibly interpret the attacks in terms of size, frequency, and location. The intent of using attack strategies rather than specific attack scenarios was to fully capture the range of possible attack alternatives Al Qaeda leadership might be considering.

2.4 Decision Tree Probability Estimates

A critical component to modeling the desirability of attack strategies is the uncertainty that Al Qaeda leaders might have regarding their success in executing each of the attack types. Research has shown that a terrorist attack operates much like any other complex business project, starting with an attack planning phase, followed by the actual preparations for the attack and culminating with attack execution [Rosoff and Winterfeldt, 2007]. Figure 2 is a simple event tree illustrating these three critical phases; event nodes to the right of an arc are conditional on the preceding events (to the left).

In the planning phase, the acquisition of material is instrumental to ensuring that the attack strategy is viable. Avoiding interdiction by anti-terror forces is an ongoing concern for the terrorist leader during the preparing phase, including such intermediary tasks as bomb building and casing of the target, both essential to attack success. Lastly, the execution phase refers to the critical final steps involved

in whether or not the attack will be carried out successfully. Here emphasis is on whether the triggering device is effective or if the executioner carries out the attack. If at any point within the planning, preparing or execution phases a task is not successfully completed, the attack is assumed terminated.

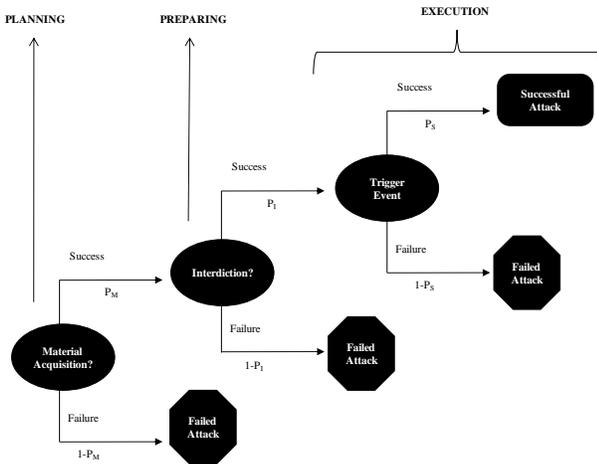


Figure 2. Attack Event Tree

Each of the planning, preparing and execution phases is associated with a probability of detection and disruption of the attack. To determine how these probability estimates affect the overall attack success probability, the proxy terrorists estimated the probability of success of each phase for each attack type. Each probability estimate varied depending upon the difficulty of the task as perceived by the proxy.

The elicitation process was two-fold [Hora, 2007]. First, proxies provided a preliminary estimate of the probability of success for a given attack strategy of obtaining material, successfully avoiding interdiction, and execution of attack. Next, they considered the uncertainty in their estimates. To account for proxy uncertainty, general beta distributions were specified for each probability estimate. Proxies specified an interquartile range (25th and 75th percentiles on the cumulative distribution) for each probability estimate. Beta parameters were estimated using the @Risk software (Best Fit module) for each distribution using these 2 points on the cumulative beta distribution ($F = .25$ and $F = .75$.)

The resulting probability distribution was then compared to the probability originally assessed and proxies were allowed to adjust their estimates to resolve inconsistencies. For example, one expert cited that the probability of a terrorist successfully acquiring radioactive material might vary depending on the number of persons involved and the type of material acquired. Another suggested that the probability of successful execution might vary depending on whether the attack was carried out using a triggering device or a suicide bomber.

2.5 Attribute Definition and Measurement

An important part of multi-attribute utility function assessment is specifying attributes that allow for comparison of the alternative attack strategies relative to the fundamental objectives defined above. Proxies provided both measures, units by which the attributes would be defined, as well as scales indicating the range over which the measures are defined. Table 2 shows how the attribute units and scales were defined by one proxy. When thinking like a terrorist, he found it easiest to base his units and scales off his understanding of outcomes following the attacks of September 11th (9/11).

ATTRIBUTE	MEASURE (UNIT)	SCALE
Short term economic	% of 9/11	0-400%
Long term economic	% of 9/11	0-500%
Recruitment	% of 9/11 militaristic recruitment	0-400%
Popular support (for Al Qaeda)	% of 9/11 popular support	0-170%
Damage to Al Qaeda	% of 9/11 organization damaged	0-60%
Instill fear in U.S.	% life change (relative to 9/11)	0-80%
Worldwide U.S. support (post attack)	% of 9/11 worldwide support	0-100%
Kill Americans	# killed	0-100,000
Funding	% of 9/11 annual funding	0-400%
Cost	Dollars (relative to 9/11)	0-\$450,000
Resources	# of people required (relative to 9/11)	0-25

Table 2. Attribute Units and Scales Defined for One Proxy

In developing the model, proxies were encouraged to define the attribute scales as they desired (as opposed to forcing a consistent set of units and scales across proxies). This enabled the proxies to be comfortable with their choices and justification of input.

Given the limited number of international terrorist attacks upon the U.S. and the inability to interview Al Qaeda leaders directly there is a level of uncertainty surrounding the proxy's interpretation of the terrorist leader, as well as the analysts' interpretation of the proxy's perspective. As a result, proxies provided uncertainty distributions over the impact of the attributes relative to each attack type. Uncertainty scales were defined by each proxy using general beta distributions on the entire matrix of 9 attack strategies by 11 attributes. For each of the 99 cells of the matrix, each proxy provided an estimate of the score obtained for the particular attack strategy on the particular scale.

The initial assessment was conducted assuming a successful attack, but modified assessments were also obtained for the three unsuccessful end nodes of the event tree described previously in Figure 2. As with the attack event tree in Section 2.4, the uncertainty distribution elicitation process was two-fold. First, the estimate was assessed as a median, such that the proxy indicated 50-50 chance that the true score was above or below the estimate.

Then, the beta distribution was obtained by assessing both a range (minimum and maximum possible scores) and an interquartile range (25th and 75th percentiles on the cumulative distribution). The Best fit module of @Risk was again used to calculate parameters of the beta distribution consistent with these four estimates. The median of the beta distribution fit to these estimates was then compared to the median obtained in the original estimate, and the proxy was allowed to make adjustments to resolve inconsistencies. The proxies were provided with several fractiles of the resulting distributions, including more extreme percentiles (5th and 95th), in addition to the values corresponding to the matching points (25th, 50th, and 75th).

The resulting beta distributions captured both the degree of overall uncertainty in the proxy estimates, as well as any skew in the direction of the uncertainty. The produced distribution is intended to quantify uncertainty about how well a given measure meets what is perceived.

By having proxies characterize the attributes using different scales, ranges and uncertainty distributions, the uniqueness of the model as a tool for capturing perceived Al Qaeda leader motivations and beliefs is best preserved. In choosing such a result, researchers were unable to compare attribute values across proxies. However, given the nature of utility models, we still were able to compare relative expected utilities and probability estimates produced as the model final output.

As noted above, the initial scores provided for each proxy's attribute matrix were constructed conditional on a successful attack. When accounting for attack feasibility, there are instances where the attack fails – whether it is during material acquisition, interdiction during transport, or unsuccessful execution. For unsuccessful attack outcomes, many attributes will have quite different scores. If, for instance, the terrorist is unable to acquire the radioactive material for a dirty bomb, certain costs and resources still are expended in the process. However, only a fraction of the originally estimated costs and resources for the full attack is tapped. There also would be “no impact” on many attributes, resulting in a score on one of the scale endpoints. Table 3 displays how the utility values of these select attributes might be reconfigured in the event of a failed attack.

	Failed Material Acquisition	Attack Interdiction During Preparing	No impact (utility = 1)
Minimize blowback to Al Qaeda	No impact (utility = 1)	% of execution utility	No impact (utility = 1)
Minimize U.S. worldwide support	No impact (utility = 1)	% of execution utility	No impact (utility = 1)
Attack cost	% of execution utility	full execution utility	full execution utility
Attack resources	% of execution utility	full execution utility	full execution utility

Execution utility refers to the utility value assigned to the attribute assuming a successful attack.

Table 3. Impact of “No Attack” on Attributes by Alternative’s Score Matrix

2.6 Risk Attitudes Across Attributes

The first instance of variation in proxy perceptions of Al Qaeda leaders’ motivations and beliefs was exemplified through attribute definition and uncertainty about impacts of attack outcomes on the scores for attack strategies on each attribute. A second variation across proxies is the single attribute utility functions that capture the terrorist leader’s attitude toward risk. While the proxies may share similar attack attributes, their perspectives toward these are not always in agreement. Each proxy can have different attitudes toward risk and in turn, may be willing to accept different levels of risk.

Utility functions are the measurement tool traditionally used to capture an individual’s attitude toward risk. The direction of the utility function indicates whether that individual is risk averse, neutral, or seeking. Through the acquisition of certainty equivalents from the proxies, we were able to estimate the nature of their risk attitudes toward each attribute. To accomplish this, an exponential utility function was estimated for each proxy for each of the 11 attributes by assessing a certainty equivalent on each attribute for a 50-50 gamble between the worst and best outcomes for each attribute. The proxy was asked to estimate a sure outcome that would make the terrorist leader indifferent between playing the gamble and taking the sure thing. Figure 3 displays two proxy’s decision trees capturing the certainty equivalent for the long term economic impact attribute.

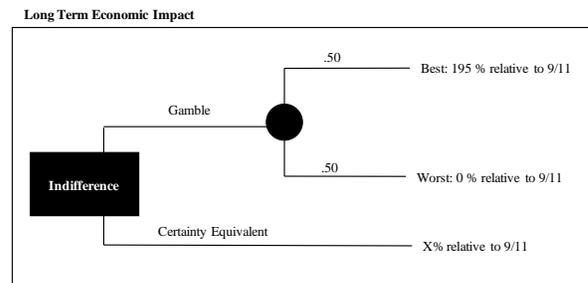


Figure 3. Long Term Economic Impact Certainty Equivalent

Proxy 1 evaluated long term economic impact in terms of damage relative to 9/11. Proxy 1 perceives 55% of the economic impact following 9/11 to be of equal value to the gamble between his best (195%) and worst (0%) estimates relative to 9/11. Proxy 2 evaluated long term economic impact in terms of a percentage of \$2 trillion. Proxy 2 perceived a terrorist leader would value 30% of 2 trillion as much as the best-worst gamble. Interestingly, while the proxies used different units for attribute definition, their certainty equivalents represented a similar percentage of the total measure, roughly 30%.

As with the definition of attribute uncertainties across proxies, generalized beta distributions were assessed for the proxy’s uncertainty about the terrorist leader’s certainty equivalent for the gamble [Hora, 2007]. (The same 4-point elicitation procedure of obtaining a minimum, maximum,

and an interquartile range was used as for the score matrix. An iterative process for comparing the median of the obtained distribution to the original estimate was also used to resolve inconsistencies.)

2.6 Value Tradeoffs Across Attributes

Another essential component of the MAU model is the proxies' perception of how terrorist leaders assess the tradeoffs among the attack attributes. Each proxy rank ordered the attack attributes using swing weights [Keeney et al, 1990]. They assigned one attribute for which the change (swing) from worst to best represented the largest impact for the terrorist leader in terms of the overall objective – committing a terror attack. All remaining attributes were assigned a percentage between 0 and 100% to reflect relative desirability of changing (swinging) a score from worst to best on that attribute.

Generalized beta distributions were assessed over each assigned swing weight to reflect uncertainty about terrorist leaders' relative value of the swing from worst to best [Hora, 2007]. The same 4-point elicitation procedure of obtaining a minimum, maximum, and an interquartile range was used as for the tradeoff uncertainty matrix.

Lastly, the swing weights were normalized to sum to 1.0. Table 4 shows the mean normalized swing weight assessments for all four proxies. While two of the proxies prioritized logistical feasibility, such as cost of the attack, one other emphasized economic impacts, and the fourth stressed the psychological implications of an attack – the ability to instill fear.

	ST Economic Impact	LT Economic Impact	Recruitment	Popular Support	Retaliation	Instill Fear
Proxy 1	.15	.16	.08	.10	.04	.12
Proxy 2	.10	.14	.07	.10	.05	.10
Proxy 3	.05	.05	.07	.08	.08	.07
Proxy 4	.15	.14	.04	.07	.08	.16

	U.S. Support	Kill Americans	Funding	Cost	Resources
Proxy 1	.08	.12	.08	.05	.05
Proxy 2	.09	.04	.00	.24	.07
Proxy 3	.09	.08	.07	.29	.06
Proxy 4	.11	.15	.02	.03	.05

Table 4. Normalized Swing Weight Assessments

3 Results

3.1 Attack Alternative Ranking

The model includes full implementation of the event tree, with 4 possible outcomes (3 failures and 1 success), as well as the additive multi-attribute utility function, utilizing exponential single attribute utility functions and trade-offs defined by swing weights. Uncertainty is captured in all proxy assessments, including event tree probabilities, the attributes by alternatives score matrix, swing weights, and certainty equivalents. A risk profile for each attack mode strategy is generated using Monte Carlo simulation. Expected utilities are calculated as the means of these

distributions, obtained from 10,000 iterations using the @Risk software (Palisades, Inc.)

The two proxy's utility distributions in Figure 4 represent risk profiles for the no attack and explosions on mass transportation alternatives. As illustrated by the risk profile ranges and curve shapes, there was some uncertainty around the desirability of the two alternatives for the terrorist leader. The no attack utility ranges from .54 - .61, indicating the alternative is moderately desirable (relative to a 0-1 scale). Plus, the areas around the mean (.57) are equally distributed, suggesting there is limited uncertainty and variability as to the alternative's desirability within this range. Comparatively, a terrorist's desirability for an attack involving explosions on mass transportation was considerably less. The utility ranges from .21 - .41, and the distribution weight falls to the left (right skewed), meaning it is likely that the terrorist's desirability for this alternative falls below the mean (.34).

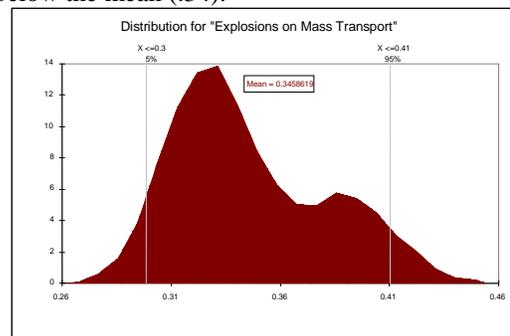


Figure 4. Utility Distribution for “Transport Explosion”

Results in Table 5 rank order the attack alternatives in terms of expected utility for each proxy. Findings indicate that when accounting for the possibility of attack failures, the attack with the highest mean expected utility was no attack for three of the four proxies (all except Proxy 2). Proxy 2's utility preference was for a smallpox attack, and this utility was only .01 greater than that for no attack. Interestingly, the utility outputs for Proxy 2 are all very similar. This suggests he might not feel that Al Qaeda leaders' preferences for attack type vary, or alternatively he has considerable uncertainty about terrorist leaders' preferences.

Attack Alternative	Proxy 1	Proxy 2	Proxy 3	Proxy 4
No attack	0.18	0.49	0.58	0.41
IED	0.15	0.45	0.34	0.21
Dam Explosion	0.14	0.47	0.29	0.15
MANPADS	0.14	0.46	0.38	0.33
Portable Nuclear Device	0.06	0.44	0.10	0.14
Transportation System	0.14	0.48	0.35	0.36
Anthrax	0.16	0.45	0.32	0.37
Dirty Bomb	0.14	0.42	0.22	0.27
Smallpox	0.09	0.51	0.17	0.14

Table 5. Attack Alternatives Expected Utilities by Proxy

Cumulative utility distributions are presented for two proxies for all nine attack strategies in Figure 5. Each curve represents a different attack strategy. Utility increases from left to right, so curves on the right generally reflect the more desired attack strategies than curves on the left.

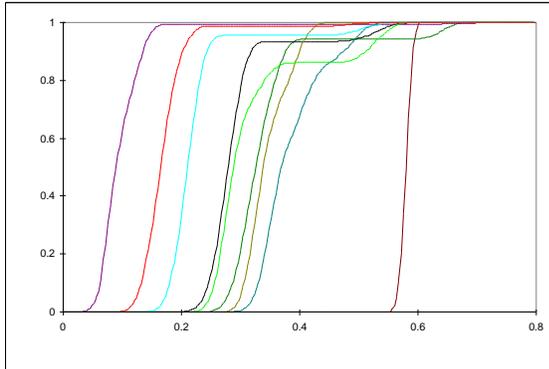


Figure 5. CDF Utility Distributions

The diagrams depict the considerable variability across attack utility and the uncertainty over that utility. For both proxies, the distributions intersect, meaning there is no stochastic dominance (one attack utility distribution preferred over the others). Also, the variations in curve shape suggest there is more uncertainty associated with some attacks compared to others. For example, the greater ‘S’ formation in proxy 1’s no attack (brown line) distribution compared to that of the smallpox attack (purple line) implies there is greater uncertainty over the utility of the smallpox attack.

Ultimately, the model produces estimated attack probabilities. These estimates were derived by sampling from the expected utility distributions. We used the risk profiles to calculate the probability that the utility for each alternative attack strategy is the maximum. Table 6 displays each proxy’s estimated subjective probabilities and specifies the most attractive attack strategy for each.

Attack Alternative	Proxy 1	Proxy 2	Proxy 3	Proxy 4
No attack	0.69	0.13	0.94	0.10
IED	0.01	0.01	0.05	0.00
Dam Explosion	0.00	0.03	0.00	0.01
MANPADS	0.00	0.09	0.00	0.24
Portable Nuclear Device	0.03	0.12	0.01	0.09
Transportation System	0.00	0.16	0.00	0.12
Anthrax	0.08	0.05	0.00	0.22
Dirty Bomb	0.14	0.07	0.00	0.17
Smallpox	0.05	0.34	0.00	0.05

Table 6. Probability of Attack

Given the possibility of not acquiring necessary material, getting caught, and failing to successfully execute an attack, the most probable attack for two of the proxies was “No attack”. Of the remaining two proxies, proxy 2 was partial to a smallpox attack, and proxy 4 to a MANPADS attack. The probability estimates for Proxies 2 and 4 are considerable lower than those for Proxies 1 and 3.

This suggests that Proxies 2 and 4 believe the attack strategies in this study were not strongly favored by terrorist leaders. However, since an attack strategy was preferred to the “no attack” alternative, it is possible the proxies might better capture a terrorist leader’s preferences. Lastly, the probability estimates across some of the attack strategies are extremely low for all four proxies, such as an IED attack and dam failure. This consensus brings to question whether these threats might now be excluded from consideration.

4 Conclusions

The proxy utility modeling approach introduced in this paper is comprised of several tasks. For starters, analysts worked closely with proxies to formulate an understanding of how terrorist leader’s values and beliefs influence their attack strategy preferences. Data was collected from proxies on attack strategy attributes, their risk preferences for the attributes, their value tradeoffs across attributes, and their estimates of attack strategy feasibility. Probability distributions also were placed over all model parameters to account for the uncertainty that the terrorist has about the attack strategy alternatives, as well as the uncertainty of the analyst’s assumption relative to the terrorist’s preferences. Analysts then used the information collected from proxies to estimate the threat of the assessed attack strategies, in terms of expected utility and probability estimates. When all pieced together, the model evaluates how perceived values and preferences intersect with perceived alternative feasibility to produce an overall probability of attack strategy selection.

Results indicate that after taking into account the possibility of not acquiring the necessary material for an attack, getting caught, or not successfully executing an attack, the attack with the highest mean expected utility was “No attack” for three of the four proxies. The remaining proxy’s utility preference was for a smallpox attack. When the probability estimates were calculated two of the three proxies were consistent in favoring the “No attack” alternative. The proxy that favored the smallpox attack preferences also did not waiver. However, the fourth proxy’s output showed a preference switch from “No attack” to a MANPADS attack.

4.1 Model Challenges

By design, the proxy value model is unable to account for changes in the objective and attribute inputs over time. Some of the variability in proxy terrorist leader values will be captured by the uncertainty distributions within the model. However, significant political, economic, and social changes will likely occur, resulting in the need to restructure and redefine some of the core inputs of the fundamental objectives hierarchy. For instance, a terror organization’s leadership may change. Whether it is because their leadership structure has evolved, or a leader is killed or

captured, new leaders rise in the ranks. The new leaders' beliefs and motivations might diverge from that their predecessor. Alternatively, as DHS counterterrorism efforts are introduced, the terrorist leader's attack preferences and strategies might be altered to adjust to the new environment. Whatever the reason, both terrorist organizations and DHS values and preferences are bound to change, and the associated utility model should be adjusted accordingly.

This might be accomplished by reevaluating the utility model on a regular basis. However, ongoing assessments would require resources and commitment on the part of the proxies and trained analysts. To avoid some of the complications with frequent data elicitation, the development of a dynamic utility/feedback model of a terrorist organization that accounts for the fluctuations in terrorist leader beliefs and motivations might have greater long term returns.

4.2 Model Applications

While the proxy utility model's outputs are mostly illustrative of the methodology used, they do show how the model might help the Department of Homeland Security (DHS) make better decisions when responding to the terrorist threat. In any decision context, it is difficult to determine how to allocate security resources when one does not understand the nature of the threat, vulnerability and consequences. The proxy value model presents a formalized approach for understanding the influence of various terrorist motivations and capabilities on the selection of potential attack strategies.

The comparative nature of the decision model output might furthermore help DHS address some of the challenges and complexity associated with the allocation of the department's annual budget. Taking time to consider the uncertainties of terrorist leader objectives and the resulting impact this might have on alternative selection (whether the alternative is an attack type or target, or other decision problem) could help to organize the scope of the department's knowledge base. As such, by initially investing in value modeling, the remaining DHS funds can be better directed toward making our country and citizens safer from the risk of terrorism.

Lastly, DHS also might explore the contributions of the proxy value model to counter intelligence analyses. The model described in this paper focuses on using the evaluation of proxy terrorist leaders' values and beliefs. Intelligence analysts present a completely different sources of data input for the model. Intelligence analysts are trained to counter to the terrorism threat (as opposed to assess it from the perspective of the terrorist), and have access to more proprietary and recent information about terrorist activities and communications. Adjustment to the model to account for different data source could provide valuable insight and/or an interesting comparative perspective on the terrorism threat.

References

- [Crenshaw, 1995] Edited by Martha Crenshaw. *Terrorism in Context*. Penn State Press, 1995.
- [Gordon *et al*, 2005] Gordon, P., Moore, J.E., Pan, Q., & Richardson, H. (2005). *Los Angeles–Long Beach seaports radioactive plume economic impact*. Draft Report, Center for Risk and Economic Analysis of Terrorist Events, University of Southern California, 2005.
- [Hora, 2007] Steve Hora, Eliciting probabilities from experts. In *Advances in Decision Analysis: From Foundations to Applications*. Edited by Ward Edwards, Ralph F. Miles and Detlof von Winterfeldt, pages 129-153. Cambridge University Press, New York, 2007.
- [Keeney *et al*, 1990] Ralph L. Keeney, Detlof von Winterfeldt, and Thomas Eppel. Eliciting public values for complex policy decisions. *Management Science*, 36(9): 1011-1030, September 1990.
- [Keeney, 1992] Ralph Keeney, Value focused thinking. Cambridge, Ma.: Harvard University Press.
- [Keeney, 2007] Ralph L. Keeney. Modeling values for anti-terrorism analysis. *Risk Analysis*, 27(3), 2007.
- [Keeney, 2007] Ralph L. Keeney. Developing objectives and attributes. In *Advances in Decision Analysis: From Foundations to Applications*. Edited by Ward Edwards, Ralph F. Miles and Detlof von Winterfeldt, pages 104-128. Cambridge University Press, New York, 2007.
- [Keeney and Raiffa, 1976] Ralph L. Keeney, and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. John Wiley & Sons, Inc., New York, 1976.
- [Rosoff and Winterfeldt, 2007] Heather Rosoff and Detlof von Winterfeldt. A Risk and economic analysis of dirty bomb attacks on the ports of Los Angeles and Long Beach. *Risk Analysis*, 23(3), 2007.
- [Simonoff *et al*, 2007] Jeffrey Simonoff, Carlos Restrepo, Rae Zimmerman, and Zvia Naphtali. Analysis of electrical power and oil and gas pipeline failures. In *IFIP International Federation for Information Processing*, pages 381-39. Springer Boston, MA, 2007.
- [Sprinak, 2000] Ehud Sprinzak. Rational fanatics. *Foreign Policy*, 120:66-73, Sep. - Oct., 2000.
- [Victoroff, 2005] Jeff Victoroff. The mind of the terrorist: A review and critique of psychological approaches. *Journal of Conflict Resolution*, 49(1): 3-42, 2005.