SAND2006-4259C

Evaluating Risk from Acts of Terrorism with Belief and Fuzzy Sets

John Darby (jldarby@sandia.gov)
Sandia National Laboratories

ABSTRACT

Risk consists of the likelihood of an event combined with the consequence of that event. There is uncertainty associated with an estimate of risk for an event that may happen in the future. For random, "dumb" events, such as an earthquake, this uncertainty is aleatory (stochastic) in nature and can be addressed with the probability measure of uncertainty.

A terrorist act is not a random event; it is an intentional act by a thinking malevolent adversary. Much of the uncertainty in estimating the risk of a terrorist act is epistemic (state of knowledge); the adversary knows what acts will be attempted, but we as a defender have incomplete knowledge to know those acts with certainty.

To capture the epistemic uncertainty in evaluating the risk from acts of terrorism, we have applied the belief/plausibility measure of uncertainty from the Dempster/Shafer Theory of Evidence. Also, to address how we as a defender evaluate the selection of scenarios by an adversary, we have applied approximate reasoning with fuzzy sets. We have developed software to perform these evaluations.

INTRODUCTION

First, we summarize how risk is typically evaluated for a random event. Then, we briefly discuss the belief/plausibility measure of uncertainty and fuzzy sets. Finally, we discuss how risk from potential acts of terrorism can be evaluated using belief/plausibility and fuzzy sets.

RISK FOR A RANDOM EVENT

For a safety analysis, risk is concerned with random failures, such as an earthquake, and risk can be defined as the product:

$$Risk = f \times P \times C \qquad \text{(Eqn. 1)}$$

where f is the frequency of the initiating event (e.g., an earthquake), P is the response of system of concern (e.g., the fragility of a building), and C is the consequence if the system fails (e.g., the number of people killed). Note that the initiating event is expressed as a frequency, f. P is conditional on the initiating event (e.g., the magnitude of the earthquake), and C is conditional on system failure. Using Equation 1, Risk has units of consequence per year.

Typically, more than one initiating event is of concern, and there is risk from each of "i" initiating events:

$$Risk_i = f_i \times P_i \times C_i \qquad \text{(Eqn. 2)}$$

The total risk can be expressed as the sum of the risk from each initiating event:

$$Total\ Risk = \sum_i Risk_i \qquad \text{(Eqn. 3)}$$

Each initiating event is associated with a scenario, S_i, where a scenario is the combination of the initiating event and the system response. The frequency of consequence for scenario "i" is defined as $F_i \equiv f_i \times P_i$. F_i is the frequency at which consequence C_i occurs. Equation 2 expresses risk as a product. To provide more information, we can define the risk from each scenario as a risk triplet: [Kaplan, Risk]

$$Risk_i = < S_i, F_i, C_i > \qquad \text{(Eqn. 4)}$$

With this formulation we can distinguish among scenarios that have similar "Risk" as defined using Equation 2, but that have significantly different frequencies and consequences. Total Risk is the set of all risk triplets:

$$Total\ Risk = \{< S_i, F_i, C_i > over\ all\ i\} \qquad \text{(Eqn. 5)}$$

Using the risk triplet approach, for each scenario we have two values: the frequency of the consequence and the consequence given the scenario occurs. We can calculate the "exceedance frequency of consequence" for the collection of scenarios. Define $P_i(C_j)$ as the probability that consequence C_j is exceeded given scenario S_i. In general, for "i" scenarios and "j" consequences:

$$Freq(C_j) = \sum_i F_i \ x \ P_i(C_j) \qquad \text{(Eqn. 6)}$$

where $Freq(C_j)$ is the frequency of exceedance of consequence value C_j. Uncertainty in F_i and C_i is expressed using a probability measure. For analysis of random events, probability is an appropriate measure of uncertainty.[1] The variables are random variables with probability distributions.

Consider a single scenario with uncertainty for each of the variables f, P, and C. The $\times$ operation in Equation 2 represents convolution of probability distributions under multiplication. For example, assume that based on the data available, f is modeled with a lognormal probability distribution with mean 1 x 10^{-3} per year and standard deviation 3 x 10^{-4} per year.[2] P is modeled with a lognormal probability distribution with mean 0.03 and standard deviation 0.01. C is modeled with a uniform probability distribution with minimum 1000 and maximum 7000 (mean 4000). Using equation 2, the expected value (mean) of Risk is 0.12 deaths per year.

[1] The name probability is used for two different concepts. The term P is a probability in the classical, or objective, sense; the number of times an event occurs divided by the number of trials in the limit as the number of trials is infinite. The uncertainty in P (due to insufficient information to calculate the classical probability) is probability in the subjective or Bayesian sense, and it represents our state of knowledge about the likelihood of the value P. Both concepts obey the Kolmogorov axioms that mathematically define a probability measure.

[2] There are many probability distributions available to model the uncertainty for a random variable, including: normal, lognormal, exponential, triangular, and normal. Data and expertise are required to select the appropriate probability distributions for the variables of interest.

There is uncertainty in the Risk as represented by the probability distribution shown in Figure 1.[3]

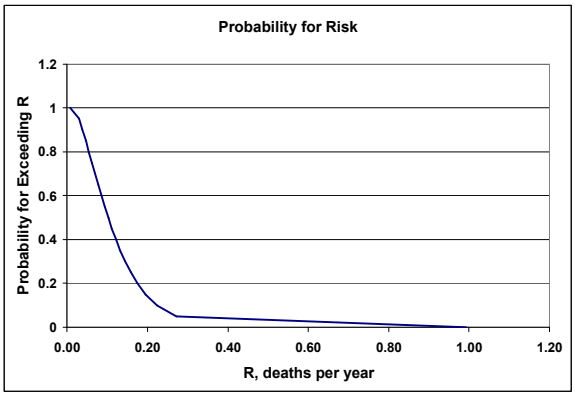

Figure 1. Complementary Cumulative Probability Distribution for Risk

For example, the probability is 95% that the risk is greater than 0.03 deaths per year. The probability is 50% that the risk is greater than 0.10 deaths per year. The probability is 5% that the risk is greater than 0.27 deaths per year.

The risk for this scenario can also be expressed as an exceedance frequency of consequence, as shown in Figure 2.[4] Due to uncertainty, there is a family of curves for selected percentiles of probability. With 50% probability the frequency of more than 5000 deaths is not larger than about 1 x 10^{-5} per year. With 95% probability the frequency of more than 5000 deaths is not larger than about 2 x 10^{-5} per year.

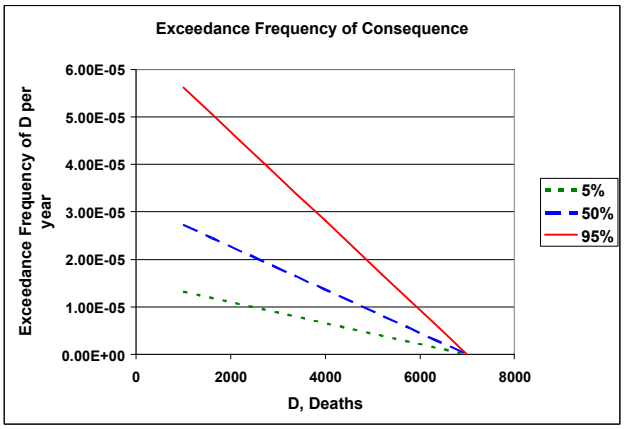

Figure 2. Exceedance Frequency of Consequence

UNCERTAINTY

Over the last 50 years, mathematicians and logicians have developed measures of uncertainty that are more general than probability, and that specifically address epistemic uncertainty. The general references provide details.

The belief/plausibility measure of uncertainty from the Dempster/Shafer Theory of Evidence is an extension of the probability measure of uncertainty that can better capture epistemic uncertainty. Belief/plausibility is a superset of probability and under certain conditions belief and plausibility both become probability. Under other conditions belief/plausibility become necessity/possibility, respectively.[5] Belief/plausibility addresses a type of uncertainty called ambiguity. Ambiguity is uncertainty associated with predicting an event in the *future*.

A simple example illustrates the difference between aleatory and epistemic uncertainty, and the use of a belief/plausibility measure. If I have a fair coin, heads on one side tails on the other with each side equally likely, my uncertainty as to the outcome of a toss- heads or tails- is aleatory. The probability of heads is ½ and the probability of tails is ½. My uncertainty is due to the randomness of the toss. Suppose however that I do not know the coin is fair; the coin could be biased to come up heads, or the coin could even be two-headed or two-tailed. Now I have epistemic uncertainty; my state of knowledge is insufficient to assign a probability to heads or tails, all I can say is the likelihood of heads (or tails) is somewhere between 0 and 1. To consider epistemic uncertainty as well as aleatory uncertainty, a superset of probability called belief/plausibility can be used as the measure of uncertainty. Using belief/plausibility, with total ignorance about the coin, the Belief that the toss will be heads is 0 and the Plausibility that the toss will be heads is 1; similarly, the Belief that the toss will be tails is 0 and the Plausibility that the toss will be tails is 1. Belief/Plausibility form an interval that can be interpreted as giving the lower and upper bound of probability. If I have specific enough information, both belief and plausibility reduce to a single value, probability. Figure 3 illustrates this concept. Note that epistemic uncertainty can be reduced with more information; if I toss the coin a few times and a heads and a tails occur, I know the coin is two sided; with more tosses I can evaluate the fairness of the coin. The aleatory uncertainty cannot be reduced with more information.

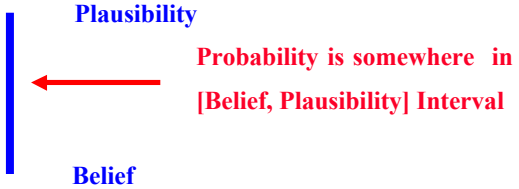

Figure 3. Belief/Plausibility as Bounds on Probability

In addition to ambiguity, we have another type of uncertainty called vagueness. We have vagueness when

[3] For the examples in this paper, the Crystal Ball software program, version 7.2, was used to convolute probability distributions under algebraic operations.

[4] The curves in Figure 2 are straight lines due to the use of a uniform probability distribution for consequence for the simple example.

[5] To be precise, if the focal elements are singletons, belief/plausibility both become probability. If the focal elements are nested, belief/plausibility become necessity/possibility, respectively.

we use linguistics (words) to classify events; for example, yesterday was "sunny", public confidence in the stock market is "high", etc. Vagueness is uncertainty as to how to classify a *known* event. For example, assume we know how tall John is, but instead of saying John is 6 feet 2 inches tall we categorize John as "tall" without a precise definition of "tall". The linguistic (word) "tall' is vague. Vagueness can be addressed using the mathematics of fuzzy sets.

A simple example of fuzzy sets is as follows. Consider a random variable for consequence as "the number of deaths from a terrorist attack" for which we take the range as $[0, 5x10^6]$. For estimating the consequence from a particular scenario we may choose to reason at a higher level than the specific number of deaths for two reasons: (a) there is too much uncertainty to distinguish between say 1000 and 2000 deaths, and (b) when comparing scenarios with widely different consequences, such as blowing up a building to detonating a nuclear device, we have orders of magnitude difference in the consequence. Suppose we partition the range with crisp sets commensurate with the "accuracy" to which we wish to measure consequence; for example, $[0, 10)$, $[10, 100)$, $[100, 1000)$, $[1000, 1 \times 10^4)$, and $[1 \times 10^4, 5 \times 10^6]$. We have defined sets, subsets of the range, at the "fidelity" to which we wish to reason. We can also assign names to these sets: "minor" for $[0, 10)$, "moderate" for $[10, 100)$, "high" for $[100, 1000)$, "major" for $[1000, 1 \times 10^4]$, and "catastrophic" for $[1 \times 10^4, 5 \times 10^6]$. We have assigned a linguistic (name) to the crisp sets of interest. But there is a problem with our crisp sets. If 999 people die the consequence is "high" but if 1000 people die the consequence is "major"; although the crisp sets solve the problem of reasoning at too fine a level, they suffer from the problem of sharp boundaries. We really want to consider 999 deaths as both high *and* major to some degree, and we can do so by making our sets fuzzy. Specifically we define "minor" as "up to *about* 10", "moderate" as " between *about* 10 and *about* 100", "high" for "between *about* 100 and *about* 1000", "major" for "between *about* 1000 and *about* 1 x 10^4", and "catastrophic" for "greater than *about* 1 x 10^4". Degrees of membership can be assigned to these fuzzy sets as indicated in Figure 4.

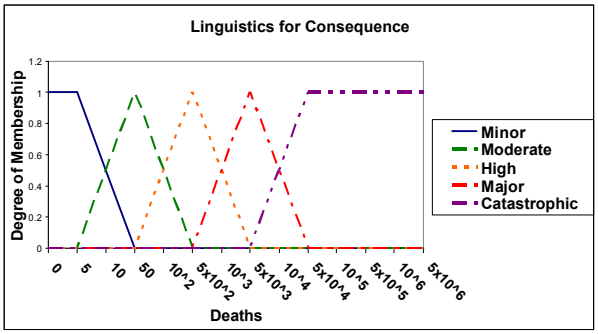

Figure 4. Fuzzy Sets for Consequence (Deaths)

Uncertainty involving both ambiguity and vagueness can be addressed by extending belief/plausibility to fuzzy sets. [Yager, 1986] Thus, we can apply the belief/plausibility measure of uncertainty to fuzzy sets.

For example, given degrees of evidence assigned to crisp intervals in the range for deaths, such as 0.7 for [10, 1000] and 0.3 for [1, 50,000], we can calculate the belief/plausibility for any fuzzy set defined for deaths.[6] An example of calculating belief/plausibility for fuzzy sets is subsequently discussed.

RISK FOR AN ACT OF TERRORISM

Risk from a terrorist act is similar in construct to Risk from a random event. For a terrorist act, we can define risk as a combination of: Threat, Vulnerability, and Consequence.

Threat is the initiating event (the terrorist act), Vulnerability is the system response (the security system response to the terrorist act), and Consequence is result of concern (e.g., deaths). Therefore, for a terrorist act, risk can be expressed in the same form as Equation 1:

$$Risk \quad = \quad f_A \times P \times C \quad \text{(Eqn. 7)}$$

where Threat is measured by f_A, the frequency of the terrorist act, Vulnerability is measured by P, the probability that the act defeats the security system in place, and C is consequence.[7]

For a terrorist act, risk is dependent on the scenario. Here, scenario is defined to include the adversary resources, attack plan, and target. Resources include attributes (equipment, weapons, number of attackers) and knowledge (perhaps from insiders). P is conditional on the scenario, since more resources raise the chance of adversary success, and C is conditional on the target in the scenario.

Risk must be evaluated for each of "i" scenarios as:

$$Risk_i \quad = \quad f_{Ai} \times P_i \times C_i \quad \text{(Eqn. 8)}$$

Equation 8 is the basic equation for an evaluation of risk from acts of terrorism.[8]

---

[6] Yager addresses the situation where the evidence is also on fuzzy sets.

[7] Sometimes risk for a terrorist act is written $Risk = P_A \, x \, P \, x \, C$ where P_A is the probability of the act. Use of P_A can cause problems if f_A is not small. P_A depends on the time of interest. Usually, the time of interest is a year. P_A can be calculated from f_A, assuming that f_A is the parameter for an exponential distribution. The probability that the scenario occurs *one or more times* within time T is $P_A(T) = 1 - \exp(-f_A T)$ which approaches 1 for large $f_A T$. It is sometimes stated that f_A in units of per year is the probability over a time period of one year; this is true only if $f_A T << 1$, since $P_A(T) \approx f_A T$ for small $f_A T$, and for T equal to 1 year, $P_A(1)$ is numerically equal to f_A. If f_A is not small, say 10 per year, the probability of the event occurring *one or more times* over a time period of a year is $P_A(1) = 1 - \exp(-10*1)$ which is approximately 1. Typically we want the number of times the consequence can occur which for large f_A is not P_A; therefore, the initiating event should be quantified as a frequency. Also, sometimes P is expressed as $(1 - P_E)$ where P_E measures the effectiveness of the security system.

[8] In practice, the P and C variables in Equation 8 can be segregated into constituent variables. For example, P can modeled as the product of two variables: (a) the probability of not detecting the gathering of resources for the scenario, and (b) the probability that the adversaries defeat the security system in place at the target. C can be the sum of many

Evaluation of Equation 8 for an intentional terrorist act is much harder than evaluation of Equation 2 for a "dumb" random event. The uncertainty associated with a terrorist act involves significant epistemic uncertainty whereas the uncertainty involved with a random event is mostly aleatory. A terrorist attack is not a random event, it involves a specific scenario that is selected, planned, and implemented by the adversary. Consider the failure of a specific building in response to an earthquake, a random event. The risk from the earthquake considers the likelihood of the earthquake, the response of the building to the earthquake, and the number of people killed if the building fails. The magnitude of the earthquake is independent of the fragility of the building. However, for an intentional terrorist attack against the building, the adversary estimates the resources required to destroy the building based on an evaluation of the fragility of the building, and decides if the potential consequences are worth the effort to bring the resources to bear necessary to destroy the building. The adversary has a choice as to which building to attack, the earthquake does not.

The terrorists have a choice, so the number of scenarios is enormous (hundreds of millions). Even if we as a defender focus on a small subset of targets for evaluation, such as Department of Defense (DoD) nuclear weapons sites, the terrorists may choose targets outside our consideration, such as Hoover Dam. A complete evaluation of f_{Ai} must address that choice.

There is significant epistemic uncertainty for the defender as to the scenario(s) that the adversary will select. The adversary has epistemic uncertainty as to the effectiveness of protective measures employed by the defender, including intelligence gathering efforts to prevent scenarios from being implemented, security systems in place to defeat an attack, and the effectiveness of measures to mitigate consequences.

We have developed an Adversary/Defender model for evaluating risk from a terrorist act. [Darby, Evaluation Terrorist Risk] [Snell, Adversary Mission Success] [Merkle, Grammar] We can use belief/plausibility together with fuzzy sets and linguistic approximate reasoning to evaluate Equation 8. The Defender part of the model solves Equation 8 numerically using belief/plausibility distributions from degrees of evidence assigned to each of the variables. The Adversary part of the model is a fuzzy set linguistic reasoning tool developed by "thinking like the adversary" and it provides information for f_{Ai}.

DEFENDER MODEL

The defender model will be explained using a simple example. Consider Risk for the scenario evaluated earlier, as summarized in Figures 1 and 2, where probability was used as the measure of uncertainty. Here, we address a scenario with the same consequence, but from an intentional act. We have significant epistemic uncertainty for the frequency of the attack, f_A. Based on the information available we assign the following evidence to f_A, where f_A is per year.

0.1 to the interval [1 x 10^{-4}, 0.1]

0.9 to the interval[1 x 10^{-3}, 0.01]

We have significant epistemic uncertainty in the expertise and knowledge of the adversary, as reflected in our assignment of the following evidence to P:

0.3 to the interval [0.1, 0.9]

0.7 to the interval[0.3, 0.5]

We assume the consequence, deaths, is as before, a uniform distribution with minimum 1000 and maximum 7000 (mean 4000).[9] Using the BeliefConvolution Java code written by the author, our result for Risk using Equation 8 is summarized in Figure 5. Our single probability curve of Figure 1 has been replaced with two curves, one for belief and the second for plausibility. These results reflect the large uncertainty in the information provided for f_A and P. Note that for this scenario, Belief is essentially zero for all values of Risk; small belief reflects assignment of evidence to intervals that are large, and in many applications belief will be small due to the large uncertainty in the information available.

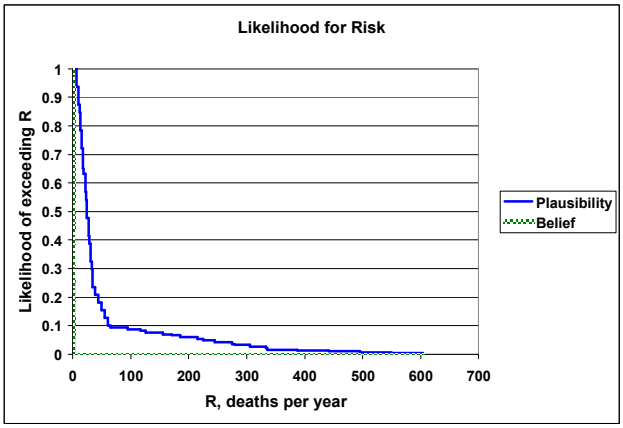

Figure 5. Complementary Cumulative Belief/Plausibility Distribution for Risk

The expected value for Risk (deaths per year) is an interval: [0.87, 47.12].

We can also summarize Risk for this scenario using fuzzy sets. We implemented Yager's technique in our BeliefConvolution code to calculate belief/plausibility for fuzzy sets. [Yager, 1986] For example, assume the fuzzy sets of Figure 4 are used, where for this example the units are deaths per year. Figure 6 expresses the results in terms of these fuzzy sets.

different types of consequences: deaths, economic loss, etc. A more detailed risk equation is discussed in a reference. [Darby, INMM]

[9] For this example, we use a probability distribution for C. In general, C may be assigned degrees of evidence over intervals and evaluated using the belief/plausibility measure.

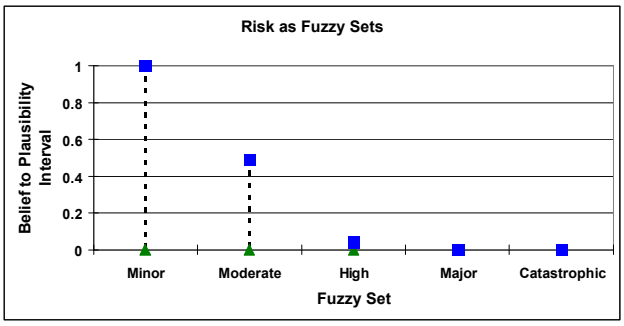

Figure 6. Risk in terms of Fuzzy Sets

We can also express risk as an exceedance frequency of consequence using a belief/plausibility measure. The $P_i(C_j)$ term in Equation 6 can be generalized to $L_i(C_j)$ where L denotes likelihood. Using belief/plausibility, $L_i(C_j)$ is an interval [Belief, Plausibility]. An upper bound for the exceedance frequency of consequence can be calculated using Plausibility for $L_i(C_j)$. [Darby, Evaluation of Terrorist Risk] For the example scenario, C has a probability distribution, so belief and plausibility for C are both probability, and L is a single value, the probability. Figure 7 summarizes the results for the example scenario.[10] Our belief is 50% that the frequency of more than 5000 deaths is not larger than about 2.0×10^{-3} per year. Our belief is 95% that the frequency of more than 5000 deaths is not larger than about 1.7×10^{-2} per year.

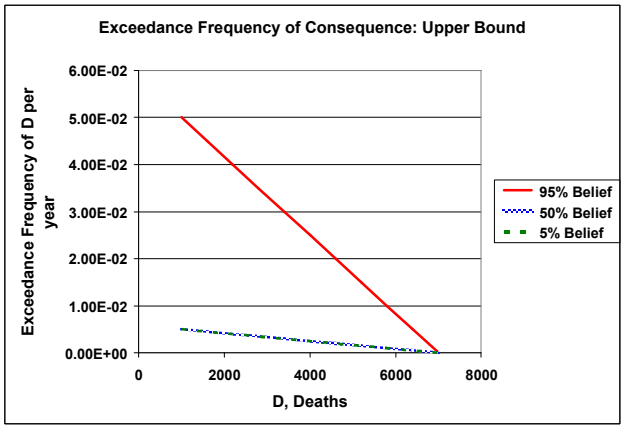

Figure 7. Exceedance Frequency of Consequence

ADVERSARY MODEL

To evaluate Equation 8 numerically using the defender model, scenarios of concern must be identified. The process of selecting scenarios requires that the defender reason from the perspective of the adversary, and this process involves a complicated consideration of many dependent factors each with significant uncertainty.

Since the adversary has a choice of scenarios, unless all the factors of importance to the adversary are "good" the adversary will discard a scenario and consider other scenarios. The adversary uses more of a "yes/no" decision process for such factors as:

1. Are the consequences of the type desired?
2. Are the potential consequences highly likely to be of sufficient magnitude?
3. Given the perceived magnitude of the consequences and the perceived level of protection, is it worth with gathering the resources needed to have a high assurance of success?
4. What are other scenarios that require fewer resources and have acceptable consequences?

That is, the adversary selects scenarios that are highly likely to succeed and maximize consequences while making effective use of resources within the constraint of the pool of resources available. The adversary spends more effort in designing the scenario for a high likelihood of success rather than estimating a precise numerical value for the likelihood of success.[11]

The adversary model evaluates scenarios using an approximate reasoning rule base for how the adversary selects a scenario. Each variable in the rule base is segregated into fuzzy sets. The fuzzy sets represent purely linguistic terms; there is no numeric definition of the fuzzy sets as in Figure 4.[12] To capture the significant uncertainty inherent in the defender thinking like the adversary, the model allows evidence to be assigned to combinations of fuzzy sets for each variable, and uncertainty is propagated up the rule base using the belief/plausibility measure of uncertainty. A Java code, LinguisticBelief, was written by the author to automate the evaluation. The adversary model is best explained by a simple example.[13]

From the perspective of the adversary, the "Expected Consequence" for a particular scenario is defined as the consequence- as perceived by the adversary- weighted by the likelihood that the scenario can be successfully accomplished- as perceived by the adversary. It is assumed that the goal of the adversary is to maximize

[10] The "percentile" curves in Figure 7 are the Belief that the frequency will not exceed the indicated value. In general,

Belief(A) = 1 – Plausibility($\bar{A}$) where A is "not A", so the Belief of "not exceeding consequence C" is one minus the Plausibility of "exceeding consequence C". In Figure 7, the curves for 50% and 5% Belief are the same, due to the interval nature of Belief/Plausibility intervals. The curves in Figure 7 are straight lines due to the use of a uniform probability distribution for consequence.

[11] That is, the adversary is not concerned with the precise likelihood of each variable of concern, such as the probability of being detected being less than 0.01. They focus on "we believe we are not likely to be detected" where not likely is ill-defined (a fuzzy set) but is understood to mean a low value (below on the order of 0.01). The decision is based on all variables of concern being acceptable to the adversary. The emphasis is on the variables of concern and how they interact rather than a precise numerical evaluation of these variables. Since the adversary has a choice, if all variables of concern are not acceptable for a particular scenario, the adversary will select another scenario.

[12] Since both the evidence and the rules are at the fuzzy set level, and we do not have the fuzzy sets defined in terms of degrees of membership, the convolution is as if the fuzzy sets were crisp. The fuzziness of the sets is considered in the assignment of evidence, not in the convolution process.

[13] The rule base is a form of approximate reasoning since it uses fuzzy sets. A simple rule base is used in this paper to illustrate the technique.

Expected Consequence. Assume the following approximate reasoning process on the part of the adversary, where "x" indicates convolution per the rule base:

- Expected Consequence = Probability Of (Adversary) Success x Consequence
- Probability Of Success = Probability Resources Required Gathered Without Detection x Probability Information Required can be Obtained x Probability Physical Security System can be Defeated
- Consequence = Deaths x Damage To National Security

Assume the following linguistics (fuzzy sets) for each variable:

- Expected Consequence = {No, Maybe, Yes}
- Probability Of Success = {Low, Medium, High}
- Consequence = {Small, Medium, Large}
- Probability Resources Required Gathered Without Detection = {Low, Medium, High}
- Probability Information Required can be Obtained = {Low, Medium, High}
- Probability Physical Security System can be Defeated = {Low, Medium, High}
- Deaths = {Minor, Moderate, Major, Catastrophic}
- Damage To National Security = {Insignificant, Significant, Very Significant}

Portions of the approximate reasoning rule base are:

Probability Of Success ***Expected Consequence*** Consequence	Low	Medium	High
Small	***No***	***No***	***No***
Medium	***No***	***No***	***Maybe***
Large	***No***	***Maybe***	***Yes***

Damage To National Security ***Consequence*** Deaths	Insignificant	Significant	Very Significant
Minor	***Small***	***Medium***	***Large***
Moderate	***Medium***	***Medium***	***Large***
Major	***Large***	***Large***	***Large***
Catastrophic	***Large***	***Large***	***Large***

Probability Physical Security System can be Defeated = High

Probability Resources Required Gathered Without Detection ***Probability Of Success*** Probability Information Required can be Obtained	Low	Medium	High
Low	***Low***	***Low***	***Low***
Medium	***Low***	***Medium***	***Medium***
High	***Low***	***Medium***	***High***

The rule base reflects the following. Expected Consequence "Yes" indicates an attractive scenario for the adversary and requires that Probability Of Success (for the adversary) be "High" and Consequence be "Large". Probability of Success "High" requires a "High" value for each of the three constituent probabilities. Consequence "Large" is from Deaths and/or Damage To National Security being severe enough from the viewpoint of the adversary.

The rule base is evaluated for each scenario of concern. Assume the following evidence assigned for a particular scenario:

- Deaths: 0.8 for {Major, Catastrophic} and 0.2 for {Moderate, Major}
- Damage To National Security: 0.1 to {Insignificant, Significant} and 0.9 to {Significant, Very Significant}
- Probability Resources Required Obtained Without Detection: 0.7 to {Medium} and 0.3 to {Medium, High}
- Probability Information Required can be Obtained: 0.15 to {Medium} and 0.85 to {Medium, High}
- Probability Physical Security System can be Defeated: 1.0 to {Medium, High}

Using the LinguisticBelief code, the following results were obtained for [Belief, Plausibility]:

- Probability of Success: [0, 0] for Low, [0.7, 1] for Medium, [0, 0.3] for High
- Consequence: [0, 0] for Small, [0, 0.1] for Medium, [0.8, 1] for Large
- Expected Consequence: [0, 0.2] for No, [0.6, 1] for Maybe, [0, 0.3] for Yes

The results can be presented graphically; Figure 8 summarizes Expected Consequence for the scenario.

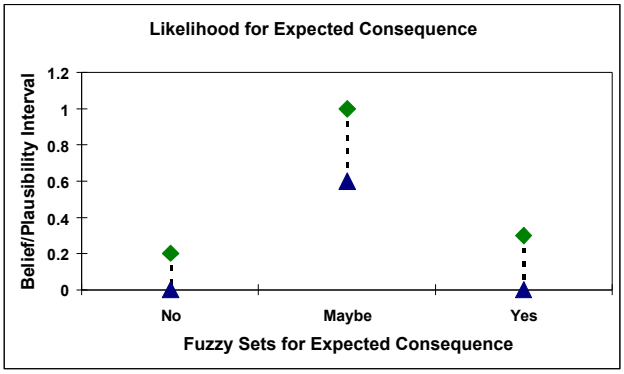

Figure 8. Expected Consequence for One Scenario

The results for this scenario indicate that although the adversary (defender thinking like the adversary) estimates a "Large" consequence to be likely (belief/plausibility of 0.8/1.0), the adversary expects Probability of Success to only be "Medium" (belief/plausibility of 0.7/1.0), resulting in an overall estimate that Expected Consequence will be "Maybe" (belief/plausibility of 0.6/1.0). Since the adversary has a choice of scenarios,

the adversary will examine other scenarios until ones with a high likelihood of "Yes" for Expected Consequence are identified.

There are many scenarios of concern. Scenarios can be ranked by decreasing expected consequence based on the plausibility for the worst fuzzy set for expected consequence: "Yes" in the prior example, sub-ranked iteratively by the plausibility of the next-worst fuzzy sets, "Maybe" and "No".[14]

SUMMARY

Evaluation of risk from acts of terrorism involves considerable epistemic uncertainty which can be captured and propagated using the belief/plausibility measure of uncertainty from the Dempster/Shafer Theory of Evidence. The risk from an act of terrorism depends on the scenario employed by the adversary and on the likelihood that the adversary selects that scenario. For a given scenario, the risk can be evaluated numerically with a defender model.

The process of selecting scenarios can be modeled linguistically with an adversary model using approximate reasoning on fuzzy sets defined for each variable. Uncertainty in the evaluation due to the defender "thinking like the adversary" is captured using the belief/plausibility measure based on evidence assigned to fuzzy sets.

ACKNOWLEDGMENTS

The use of a linguistic rule base for modeling the adversary is based on concepts from the Logic Evolved Decision (LED) methodology developed at Los Alamos National Laboratory (LANL) by Terry Bott and Steve Eisenhawer, extended to include belief/plausibility as the measure of uncertainty. The numerical model for the defender benefited from the work and suggestions of Jon Helton at Arizona State University. [Helton, Alternative Approaches] [Oberkampf and Helton]

The evaluation of belief/plausibility for fuzzy sets uses the technique developed by Ronald Yager at Iona College. [Yager, 1986]

Scott Ferson of Applied Biomathematics provided suggestions and helpful reference material during the formulation of the concepts used in this work. The RAMAS RiskCalc software (version 4.0) developed by Ferson, et al, was used to check test case results of the BeliefConvolution code.

The Laboratory-Directed Research and Development Program at Sandia National Laboratories provided support for this work. Sandia is a Multiprogram Laboratory Operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000. SAND2005-xxxx, Unclassified/Unlimited Release.

[14] As the defender thinking like the adversary, we rank by plausibility. If the actually adversary used this linguist evaluation tool to assist in the selection of scenarios, the adversary would rank by belief. This is evident in exercises conducted by members of military special forces acting as a surrogate adversary; unless they "believe" that a scenario has high certainty of success, they will discard this scenario and chose another one with less uncertainty.

SPECIFIC REFERENCES

[Darby, Evaluation Terrorist Risk] Darby, J., "Evaluation of Risk from Acts of Terrorism: The Adversary/Defender Model Using Belief and Fuzzy Sets", draft report, Sandia National Laboratories, Albuquerque, NM, to be published.

[Darby, INMM] Darby, J., "Evaluating Terrorist Selection of Attack Scenarios using Belief, Fuzzy Sets, and Approximate Reasoning", SAND2006-3385C, presented at the 47th Annual Meeting of Institute of Nuclear Materials Management, July 16-20, 2006, Nashville, TN.

[Kaplan, Risk] Kaplan, S, and Garrick, B.J., "On the Quantitative Definition of Risk", Risk Analysis, Vol. 1 No. 1, 1981 11-27.

[Merkle, Grammar] Merkle, P., "Extended Detection and Defense: 1. Adversary-Defender Modeling Grammar for Vulnerability Analysis and Threat Assessment, SAND2006-1484, March 2006.

[Snell, Adversary Mission Success] Snell, M., "Estimation of Probability of Adversary Mission Success", presented at the 47th Annual Meeting of Institute of Nuclear Materials Management, July 16-20, 2006, Nashville, TN.

[Yager, 1986] Yager, R., "Arithmetic and other operations on Dempster Shafer structures." International Journal of Man-Machine Studies, 25: 357-366, 1986.

GENERAL REFERENCES

[Dubois and Prade, Possibility Theory] Dubois, D. and Prade, H., Possibility Theory, An Approach to Computerized Processing of Uncertainty, Plenum Press English Translation, 1988.

[Helton, Alternative Approaches] Helton, J.C., Johnson, J.D., and Oberkampf, W. L., "An Exploration of alternative approaches to the representation of uncertainty in model predictions", Reliability Engineering and System Safety 85 (2004) 39-71.

[Klir and Yuan, Fuzzy Sets and Fuzzy Logic] Klir. G. J., and Yuan, B., Fuzzy Sets and Fuzzy Logic, Theory and Applications, Prentice Hall PTR, 1995.

[Oberkampf and Helton] Oberkampf, W. L., and Helton, J.C., "Evidence Theory for Engineering Applications", Chapter 10 of Engineering Design Reliability Handbook, Nikolaides, E., Ghiocel, D.M., and Singhal, S., eds., CRC Press, 2005.

[Shafer, Theory of Evidence] Shafer, G., A Mathematical Theory of Evidence, Princeton University Press, 1976.