

Probabilistic Risk Analysis and Terrorism Risk

Barry Charles Ezell,¹ Steven P. Bennett,² Detlof von Winterfeldt,³
John Sokolowski,¹ and Andrew J. Collins¹

Since the terrorist attacks of September 11, 2001, and the subsequent establishment of the U.S. Department of Homeland Security (DHS), considerable efforts have been made to estimate the risks of terrorism and the cost effectiveness of security policies to reduce these risks. DHS, industry, and the academic risk analysis communities have all invested heavily in the development of tools and approaches that can assist decisionmakers in effectively allocating limited resources across the vast array of potential investments that could mitigate risks from terrorism and other threats to the homeland. Decisionmakers demand models, analyses, and decision support that are useful for this task and based on the state of the art. Since terrorism risk analysis is new, no single method is likely to meet this challenge. In this article we explore a number of existing and potential approaches for terrorism risk analysis, focusing particularly on recent discussions regarding the applicability of probabilistic and decision analytic approaches to bioterrorism risks and the Bioterrorism Risk Assessment methodology used by the DHS and criticized by the National Academies and others.

KEY WORDS: Adaptive adversary; decision tree; event tree; terrorism risk

1. INTRODUCTION

"Probability is the guide to life."—Cicero (107 BC)

"We have to identify and prioritize risks—understanding the threat, the vulnerability and the consequence. And then we have to apply our resources in a cost-effective manner." (Michael Chertoff, Former Secretary of the Department of Homeland Security, 2006)

For more than 30 years, probabilistic risk analysis (PRA) has been a major tool for assessing risks

and informing risk management decisions by government and industry, in areas as diverse as environmental protection, industrial safety, and medical decision making. Applications of PRA to terrorism risks are new, however, and not uncontroversial. Here, we take a broad view of PRA, including any probabilistic approach involving tools like event trees, fault trees, and decision trees. We also introduce other tools such as game theoretic approaches and system dynamics, which may prove to be useful in dealing with the intelligent adversary. A major challenge in risk analysis of terrorism is the fact that terrorists, unlike nature or engineered systems, are intelligent adversaries and may adapt to our defensive measures. There has been recent criticism of PRA approaches to terrorism risk analyses, especially (but not only) by the National Research Council's Committee on Methodological Improvements to the Department of

¹ Virginia Modeling, Analysis and Simulation Center, Old Dominion University, Norfolk, VA, USA.

² Office of Risk Management and Analysis, NPPD, U.S. Department of Homeland Security, Washington, DC, USA.

³ International Institute for Applied Systems Analysis, Laxenburg, Austria.

*Address correspondence to Barry Ezell, 1030 University Blvd., Suffolk, VA 23435; bezell@odu.edu.

Homeland Security's (DHS) Biological Agent Risk Analysis (referred to hereafter as the NRC Committee). The NRC Committee has argued that because of this adaptive nature, it is problematic to assess probabilities of terrorism events or to use traditional PRA tools like event trees, suggesting alternative tools to assess the risks of terrorist events. One purpose of the article is to justify the use of PRA for terrorism risk analysis, while acknowledging its limitations. A secondary purpose of the article is to propose a pluralistic approach to terrorism risk analysis, which allows alternative approaches to be examined and tested. To this end, we examine some alternative approaches and discuss their contributions and limitations. While we do not take issue here with the possible value of these alternative approaches, we aim to make a case that (1) probabilities of terrorism events are useful to assess terrorism risks; (2) event trees can be used as part of a terrorism PRA to decompose the universe of terrorism scenarios; and (3) alternatives suggested by the NRC Committee like extended forms of games or decision trees constructed from the terrorists' perspective, like all approaches, have limitations. This article is organized in the following way. Section 1.1, provides a short background on the Department of Homeland Security (DHS) Bioterrorism Risk Assessment (BTRA) methodology and the context that motivated this article. Following the BTRA background, it concludes with a summary of the NRC Committee's criticism of the use of probability to assess the likelihood of terrorism events and the use of event trees in favor of approaches that consider terrorist events as actions that can be derived from their objectives. Section 2 details the usefulness of probabilities in bioterrorism risk analysis. Section 3 provides an overview of several tools that have been used or might be used to account for the intelligent adversary in terrorism risk. The review of tools in Section 3 is not intended to be exhaustive, but rather to note that modeling tools have limitations in dealing with the intelligent adversary. For example, some tools and approaches, while promising, may require additional development before ready for use in real-world applications while others are more mature and established. The final section summarizes, and advances, the claim that no single model or approach is sufficient to cover the entire landscape of terrorism risk and support the difficult decisions that must be made by Homeland Security decisionmakers.

1.1. 2006 Bioterrorism Risk Assessment Background

Signed in 2004, Homeland Security Presidential Directive⁴ (HSPD) 10 focused on improving the nation's ability to prevent, prepare for, respond to, and recover from terrorism attacks that employed biological agents as their means. An important component of HSPD-10 was the president's requirement for DHS to develop "periodic assessments of the evolving biological weapons threat," explaining that "the United States requires a continuous, formal process for conducting routine capabilities assessments to guide prioritization of . . . on-going investments in biodefense-related research, development, planning, and preparedness." The first national Bioterrorism Risk Assessment would be required by January 2006. To meet this requirement, in early 2005, DHS investigated three methodologies varying in complexity, depth, and community familiarity. A Technical Expert Review Panel reviewed each methodology.⁵ Based on resulting comments, and other factors, DHS determined that BTRA should primarily be a PRA-based methodology. BTRA has evolved over the years, incorporating new tools and techniques as science progresses and as program realities allow.

DHS requested the National Academy of Sciences' National Research Council (NRC) to review BTRA in 2006. The NRC's Draft Final Report, delivered in January 2008, recommended "to assess the probabilities of terrorist decisions, DHS should use elicitation techniques and decision-oriented models that explicitly recognize terrorists as intelligent adversaries who observe U.S. defensive preparations and seek to maximize achievement of their own objectives." Also, the committee chairman proposed a decision tree approach from the "terrorist point of view."⁽¹⁾

DHS identified several concerns with the NRC Committee's report.⁽²⁾ In particular, the conclusion that probability of terrorism events and event trees are not suitable for bioterrorism risk analysis appeared to be controversial and not shared by many in

⁴ Homeland Security Presidential Directive 10: Biodefense for the 21st Century, 2004. Available at www.fas.org/irp/offdocs/nspd/hspd-10.html.

⁵ Detlof von Winterfeldt was a member of the Technical Expert Review Panel in 2005.

the risk analysis community. This article challenges the NRC Committee's conclusion.

1.2. Intelligent Adversary Analysis

An essential aspect of any terrorism risk assessment is the approach used to represent and model terrorist adversaries. It is arguable that one of the best sources of information on the nature and intelligence of our adversaries, although limited, uncertain, and incomplete, is the intelligence community (IC). The IC persistently observes, collects, fuses, and assesses terrorist activities, motivations, intent, and capabilities. The ongoing challenge for DHS risk analysts, then, is how best to consult, incorporate, and transform relevant intelligence information into meaningful inputs for terrorism risk analysis, in conjunction with other models of terrorists' behavior outside of the IC.

Intelligence products exist in a range of forms, from opinions based on anecdotal information, to assessments based on tradecraft, and in other cases, technical methods and models. How, then, might DHS and the IC transform intelligence information into meaningful inputs for bioterrorism risk analysis? The NRC Committee advised DHS to model potential bioterrorists as "intelligent adversaries" as a part of its risk assessment—assuming that at each decision point in the planning of an attack, the adversary will always make the choice that maximizes his or her objectives, thus making terrorism attack probabilities *outputs* of decision models, rather than incorporating intelligence information as input.⁽¹⁾

In decision analysis terminology, the NRC Committee proposed to conceptualize the interaction between defenders and attackers in an evolving terrorist attack as a decision tree, in which the attacker's choices are modeled as decisions that maximize expected utility and the defender's choices are modeled as uncertain events, related to the relative effectiveness of the defenses. Three other possibilities are (1) a decision tree in which the defender's choices are modeled as decisions that maximize expected utility and the attacker's choices are modeled as uncertain events that are influenced by the defender's decision; (2) a decision tree in which both the attacker's and the defender's choices are modeled by decisions that maximize expected utility, e.g., an extended form of a game; and (3) an event tree that models both the attacker's choices and the defender's responses as uncertain events.

Clearly, there are advantages and disadvantages to these ways of representing attacker–defender interactions and there is no "correct" answer.

2. PROBABILITIES ARE USEFUL TO QUANTIFY THE RISK OF TERRORIST ATTACKS

In the first issue of the journal *Risk Analysis*, Kaplan and Garrick published an important paper that defined risk as the triplet of scenario, likelihood, and consequence.⁽³⁾ For the following three decades, the risk and decision analysis communities have cited this seminal paper and used many of the concepts and tools developed in it. More recently, Garrick *et al.* advocate the use of PRA for assessing terrorism risk, specifically for assessing the probabilities of terrorist attacks.⁽⁴⁾ Work based on Garcia, McGill *et al.*, Paté-Cornell and Guikema, Rosoff and von Winterfeldt, Willis, and von Winterfeldt and O'Sullivan is an example of risk analyses that use PRA, and that externally estimate probabilities of terrorist attacks as inputs.^(4–10)

Willis, McGill *et al.*, and other terrorism risk researchers operationalize terrorism risk as the product of threat, vulnerability, and consequences. More specifically, threat is usually defined as the probability of an attack (weapon, delivery mode, target, etc.), vulnerability as the probability of an attack's success given that it occurs, and consequences are the losses that occur (fatalities, injuries, direct and indirect economic impacts, among others) given a successful attack. Equation (1), then, is a common expression of homeland security risk.^(11,6)

$$Risk = P(A) \times P(S|A) \times C \quad (1)$$

Hence, a useful first-order indicator of terrorism risk is the expected consequences (loss of lives, economic losses, psychological impacts, etc.) against which the benefit of existing or potential terrorism strategies, policies, and countermeasures can be evaluated and estimated. In this probabilistic framework, the attack probabilities ($P(A)$ in Equation (1)) are for the most part agreed to be the most challenging to estimate. Quantifying $P(A)$ requires knowledge, data, or modeling about the motivations, intent, and capabilities of terrorists (largely the domain of the intelligence community), in addition to or instead of knowledge about historical attacks and their relevance to current risk.

It is very difficult to elicit absolute probability (or frequency) judgments that permit this kind of

output. However, relative judgments in terms of rank orders or ratios are easier to acquire from intelligence or other experts. For example, while it may be difficult to assess the absolute probability that a particular terrorist group will engage in a terrorism attack using nuclear materials in the United States in the next 10 years, experts can more easily reason comparatively, and might judge a “dirty bomb” attack using radiological material from a medical facility is more likely or less likely than an attack using an improvised nuclear device by considering the relative technical difficulties of executing these attacks. There is extensive literature regarding methods for eliciting uncertain probability judgments (often as probability distributions) from experts, which suggests how one might elicit probabilities in the face of intelligence complexities and uncertainties inherent in terrorism risk analysis (for a recent summary, see Bedford and Cooke, and Hora^(12,13)).

When intelligence analysts estimate a probability of attack, they are making a statement of belief about what a terrorist might do, based on available intelligence information as well as their personal experience and judgment. Apostolakis makes this crystal clear: “there is only one kind of uncertainty stemming from our lack of knowledge concerning the truth of a proposition. Distinctions between probabilities are merely for our convenience in investigating complex phenomena. Probability is always a measure of degree of belief.”⁽¹⁴⁾

There are two common arguments against the use of expert-estimated attack probabilities for terrorism risk analysis: (1) that the level of uncertainty and incompleteness associated with intelligence data prevents reasonable probability estimates from being made, even when using expert elicitation approaches that are designed to capture and represent uncertainty, and (2) that these probabilities are not static—i.e., the adversary is intelligent, observing U.S. defensive actions and shifting attack preferences accordingly.

Regarding the first argument, it is important to note that intelligence information is already in use for decision support at the highest levels in government; uncertainty and incompleteness are managed and communicating by representing judgments verbally with associated caveats. This approach, however, has historically led to some significant misunderstandings of intelligence information, a notable example being a (now declassified) 1951 National Intelligence Estimate (NIE 29–51), entitled “Probability of an Invasion of Yugoslavia in 1951.” In this

intelligence document appeared the statement: “Although it is impossible to determine which course the Kremlin is likely to adopt, we believe that an attack on Yugoslavia in 1951 should be considered a serious possibility.” When asked by State Department staff what odds the authors of the assessment placed on an attack in 1951, Sherman Kent of the National Board of Estimates replied “65 to 35 in favor of an attack.”⁽¹⁵⁾ The State Department had interpreted “serious possibility” as being “very considerably lower” than Kent’s 65/35 reply. Kent then polled the other authors of the document to determine the odds they had in mind when they agreed to the wording, observing that the odds in the minds of the authors ranged from 80/20 to 20/80 in favor of an attack.

The example above is not intended to criticize the production and communication of intelligence information; rather, it highlights an opportunity for improved clarity and understanding of uncertainty when a mathematical language for capturing and expressing degree of belief—probability theory—is used. Expression of intelligence information in a consistent manner that reflects uncertainty and is able to be incorporated into other models is helpful and arguably can improve the interpretation and utility of the information, particularly as it informs risk analysis.

Regarding the second argument against using expert-elicited attack probabilities, the adaptive nature of the adversary is certainly an important consideration. Nevertheless, it is reasonable to start with a baseline of defensive actions, current terrorist motivations, intent, and capabilities (based on data, intelligence, and other expertise), and then assess probabilities conditional on this baseline. We take it for granted that all probabilities are conditional on our current state of knowledge. While it is perhaps more difficult to spell out these conditions precisely in terrorism risk analysis, there is no fundamental difference in this type of conditioning compared to conditioning probability judgments in the case of natural or engineered systems.

Once we introduce new defensive actions, it is, of course, important and necessary to reassess these probabilities in light of the preventative, protective, or deterrence effects of the defensive actions. For example, as von Winterfeldt and O’Sullivan pointed out, the use of countermeasures to Man-Portable Air-Defense Systems (MANPADS) is assessed to have a strong deterrence effect on terrorists who may contemplate the use of MANPADS weapons (such as shoulder-fired missiles, etc.) to attack commercial

airplanes.⁽¹⁰⁾ It is important to note, however, that re-evaluation of probabilities following defensive action is not necessarily always a best estimate of risk, since for a terrorist adversary, it is next to impossible to determine whether or not, or the degree to which, the adversary is in fact (1) aware of particular defensive actions and their subsequent implications, and (2) adjusting the adversary's decisions and preferences based on awareness of defensive actions. Additional intelligence information can assist in determining the "penetrance" of U.S. defensive adjustments into the adversary's decision-making process, but any newly determined "postdefensive adjustment" risks may well be best presented to decisionmakers alongside, or in addition to, baseline risks rather than instead of them.

3. TOOLS FOR TERRORISM RISK ANALYSIS

Probabilities associated with complex events are difficult to assess directly, and it is therefore often useful to decompose these events into components and to determine the overall event probability by assembling the components' probabilities using standard probability calculus. There are many alternative decomposition tools, including event trees, fault trees, decision trees, influence diagrams, and belief nets. When the intention is to divide a very large universe of events into a structured set, event trees are useful as part of a baseline assessment of terrorism risk, beginning with an initial choice of weapon and target, and following through the path from attack, through success or failure, to eventual consequences. Event trees have been used to decompose terrorism scenarios in a number of efforts.^(16–18) Rosoff and von Winterfeldt use event trees to track the paths to failure or success of a dirty bomb attack and von Winterfeldt and O'Sullivan use a combination of decision and event trees to quantify the costs and benefits of countermeasures to MANPADS.^(8,10)

We present three categories of tools for use in PRA as it applies to terrorism risk, beginning with an introduction of logic trees under which we group forward logic trees and fault trees. Next, we briefly review additional methods—influence diagrams, systems dynamics models, and Bayesian networks (BN)—as potentially useful in transforming conceptual terrorist actions into computational models. For the final category, we discuss game theoretic approaches. For each we discuss the potential advantages and limitations.

3.1. Logic Trees

Logic trees are important tools for exploring the scenario space, analyzing uncertain events, defining scenarios, and assessing risk.⁽¹⁹⁾ The use of logic trees in probabilistic seismic hazard analysis has a long history, ranging from weighting of a few alternative assumptions to full uncertainty treatment for all of the inputs to a probabilistic assessment. Logic tree analysis consists of specifying a sequence of assessments that must be made in order to perform an analysis and then addressing the uncertainties in each of these assessments in a sequential manner. Thus, it provides a convenient approach for breaking a large, complex assessment into a sequence of smaller, simpler components that can be more easily addressed.⁽²⁰⁾ In this next section, we divide logic trees into two categories: (1) probability, event, and decision trees, and (2) fault, attack, and success trees. Where some may draw a serious distinction between probability, event, and decision trees, they fundamentally all use forward logic in their design. Parnell *et al.* arrived at a similar conclusion in a report by the Homeland Security Institute for DHS that represented "consensus among the authors" and detailed 20 risk assessment frameworks.⁽²¹⁾ In this article, we do not attempt to recount what has already been published. Instead, we narrow our focus to trees and game theory; two areas (minus PRA event trees) that the NRC Committee strongly recommended as the appropriate way to do bioterrorism risk analysis.

3.1.1. Probability, Event, and Decision Trees

Probability trees model a sequence of uncertain events in order to calculate the probabilities of events in the outcome space (Fig. 1). A probability tree is a succession of circular nodes (uncertain state variables) with branches. The branches emanating from each node represent the different possible values of the uncertain variables associated with the node. Probability trees have the following properties: (1) event \ nodes and branches; (2) forward logic; and (3) downstream events conditioned on previous nodes. Probability trees have many uses such as (1) to graphically represent the fundamentals of probability theory; (2) to describe probabilistic relationships between two or more events; and (3) to serve as the mathematical foundation for more advanced tree structures such as event trees or decision trees.

Event trees inductively model the sequences of events that lead to consequences.⁽²²⁾ Event trees

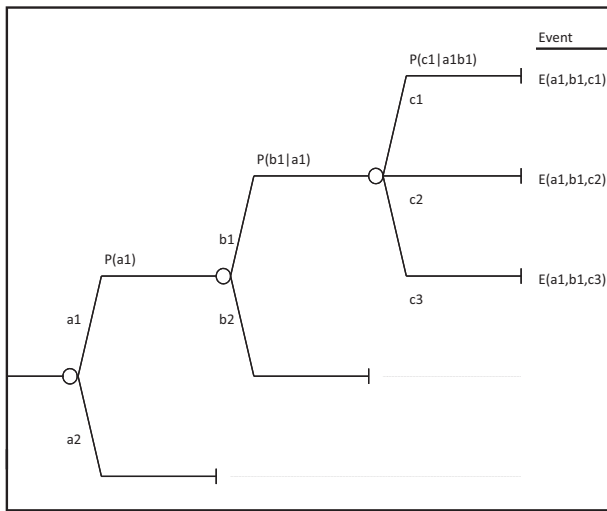


Fig. 1. Probability tree.

have the following properties: (1) events (nodes) and branches; (2) forward logic; and (3) downstream events conditioned on previous events. Event trees are an extension of probability trees by adding: initiating event, mitigating events, and consequences. Consequences are added for each probability path. Event trees have been used in many fields. For large systems, event trees have been used in nuclear reactor safety studies. Ezell *et al.* employed event trees to understand cyber risk to supervisory control and data acquisition systems for water supply.⁽¹⁶⁾ In PRA, event trees operate by identifying the likelihood of any given probability path (from initiating event through the leaves of the all tree branches).

Probabilities are assigned to event tree branches to represent the relative likelihood or degree of belief about the outcome of each branch. Probabilities at a given node are assessed conditionally on the assumption that all the branches leading to that node represent the true states of the preceding events. Because they are conditional probabilities for assumed mutually exclusive and collectively exhaustive events, the sum of the conditional probabilities at each node is unity.⁽²³⁾

Event trees have been used for very large complex systems. For example, NASA has a mature program using PRA for decision making and managing project risk—Mars missions, Space Station construction, Space Shuttle flights, etc. Compared to the BTRA event tree, NASA PRAs involve extremely large sets of unknowns. While it is desirable to create small and compact event trees that are simply described, this is often inadequate for the represen-

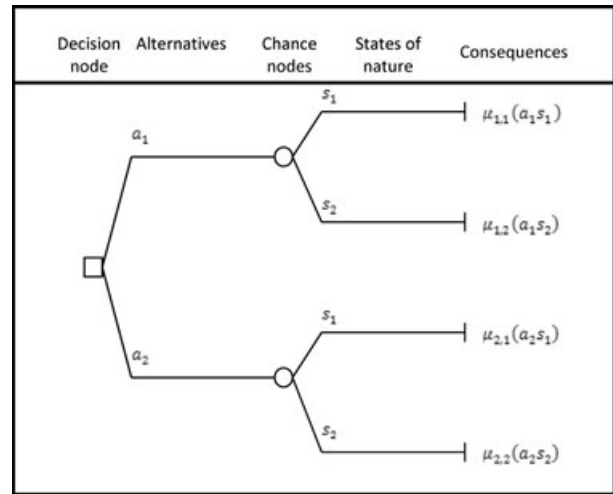


Fig. 2. Decision tree.⁽²⁶⁾

tation of real uncertainties. Consider as an example, a comparison between the BTRA and NASA PRA. The BTRA is comprised of one event tree, 16 events, and 74 branches. A NASA Space Shuttle PRA has 5,000 event trees, 6,000 events, and 2,000,000 branches, and approximately 100 off-line supporting models.⁽²⁴⁾

Decision trees are logic trees that include decision nodes in addition to events. A decision tree is effectively a diagram of a decision, read left to right.⁽²⁵⁾ The leftmost node in a decision tree is the root node and is usually a decision node (represented by a square). Branches emanating to the right from a decision node represent the set of decision alternatives that are available. Small circles in the tree are chance nodes that represent uncertainty in outcomes (Fig. 2). In the same fashion as probability trees and event trees, probabilities are assigned to the branches, referred to as states of nature, emanating from the chance nodes in the tree. The leaves of each path through the tree are called endpoints. Each endpoint represents the final outcome of a path from the root node of the decision tree to that endpoint.⁽²⁶⁾ In decision analysis, decision trees are used as a decision support tool to find the alternative with the best expected value. In the terrorism context, decision trees can structure the attacker's actions as decisions and the defender's as chance nodes vs. structuring the defender's actions as decisions and the attacker's actions as chance nodes. The NRC report emphasizes the use of decision trees that structure the attacker's actions as choices and the defender's responses as chance nodes. Other studies do the reverse.⁽¹⁰⁾ When both the attacker's and the defender's choices are

modeled as decision nodes, such approaches are considered to be in the domain of game theory.

Logic trees have limitations. First, probability assignments are usually based significantly on subjective judgments (1) because of the often limited availability of numeric data, and (2) because human judgment is often needed to weigh alternative interpretations of whatever data are available. In the context of the BTRA, intelligent adversary decisions are modeled as probability estimates, which are structured to incorporate the intelligence community expert knowledge about the intelligent adversary's preferences and capabilities. Another limitation is that terrorist choices are not random events such as failure events in engineered systems or natural hazard events. At best, the probabilities assigned by the intelligence community represent its belief as to the choices terrorist will make. Terrorists are not static and change their preferences based on what we choose to do so probability judgments must be updated as the situation changes. Event tree PRA is therefore only a "snapshot" of threat, vulnerability, and consequences for a given time period.

In BTRA, DHS used an event tree and modeled terrorist decisions as probability estimates elicited from the intelligence community—estimates that were elicited specifically to incorporate the adversary's level of sophistication, or "intelligence" from the intelligence community's perspective. The NRC Committee, however, rendered an opinion that the government should use decision trees from the terrorist's perspective, where at each decision point in the tree, an adversary will always choose the "branch" that maximizes the consequences of an attack. Decision trees (as the Committee points out) have the advantage that they are a simpler form of analysis, and more easily explainable.

Using decision trees to model terrorism from the perspective of terrorists (as proposed by the Committee) requires two important assumptions. First, it assumes ideal adversary intelligence and rationality—that the adversary *knows* which branch choice at a particular decision node best maximizes consequences at tree endpoints. For most terrorist attack paths, and certainly for technically complex attack paths requiring some degree of sophisticated knowledge and skills such as those for bioterrorism, this can be a very unsafe assumption for the U.S. government.⁽²⁷⁾ This assumption is problematic for a few reasons—first, successfully executing a WMD terrorism attack such as bioterrorism is a technically challenging enterprise, often with counterintuitive rela-

tionships between technical planning decisions and resulting consequences (e.g., terrorists could assume that a particular mode of production for "virus A" might be optimal in terms of producing casualties when used in attack since that mode is known to be optimal for "virus B"; such assumptions in biology are often false due to the complexity of biological systems). As such, many decisions that may appear to be the consequence-maximizing ones may in fact be the exact opposite, and the level of technical proficiency—an aspect of terrorist groups that is studied by the intelligence community—can be the determining factor in whether or not the adversary knows this. Further, the assumption in proposed decision tree approaches that terrorist adversaries have a perfect working knowledge of all defensive countermeasures arrayed against them is potentially problematic for two reasons; first that the adversary really does know all the details of our defenses, and second that even if the adversary did, he or she would act rationally in possession of that information. Were intelligence information to indicate that either of these were not the case, the objective functions for the terrorist would be undermined to the point that the utility of the analysis for estimating risk from overidealized terrorists would be greatly reduced.

The second critical assumption, also related to the adversary's objective function in the decision mode, is that the intelligence community (or anyone else for that matter) knows the objectives the adversary is trying to maximize in the first place. While consequence maximization may be a good guess, it is not easy to determine which types of consequences are most desirable to the adversary, and in what proportions. This is an important aspect of the analysis in that decision points for the adversary may be consequence maximizing for some consequence measures (such as fatalities), but not necessarily for others (such as economic or psychological consequences). For example, consider the 2001 anthrax attacks in the United States, in which four people were killed and 18 (known) infected. It has been noted that the widespread panic, closure of government and postal facilities, and massive public expenditures and preventive actions were of higher impact than the four fatalities, and that psychological and economic impacts may be more important adversary objectives than killing Americans.⁽²⁸⁾ However, one way to manage this challenge is to examine data, and elicit experts that seek to assess the *outputs* of terrorists' objective functions—that is, what adversaries are actually doing or planning. Specifically, while

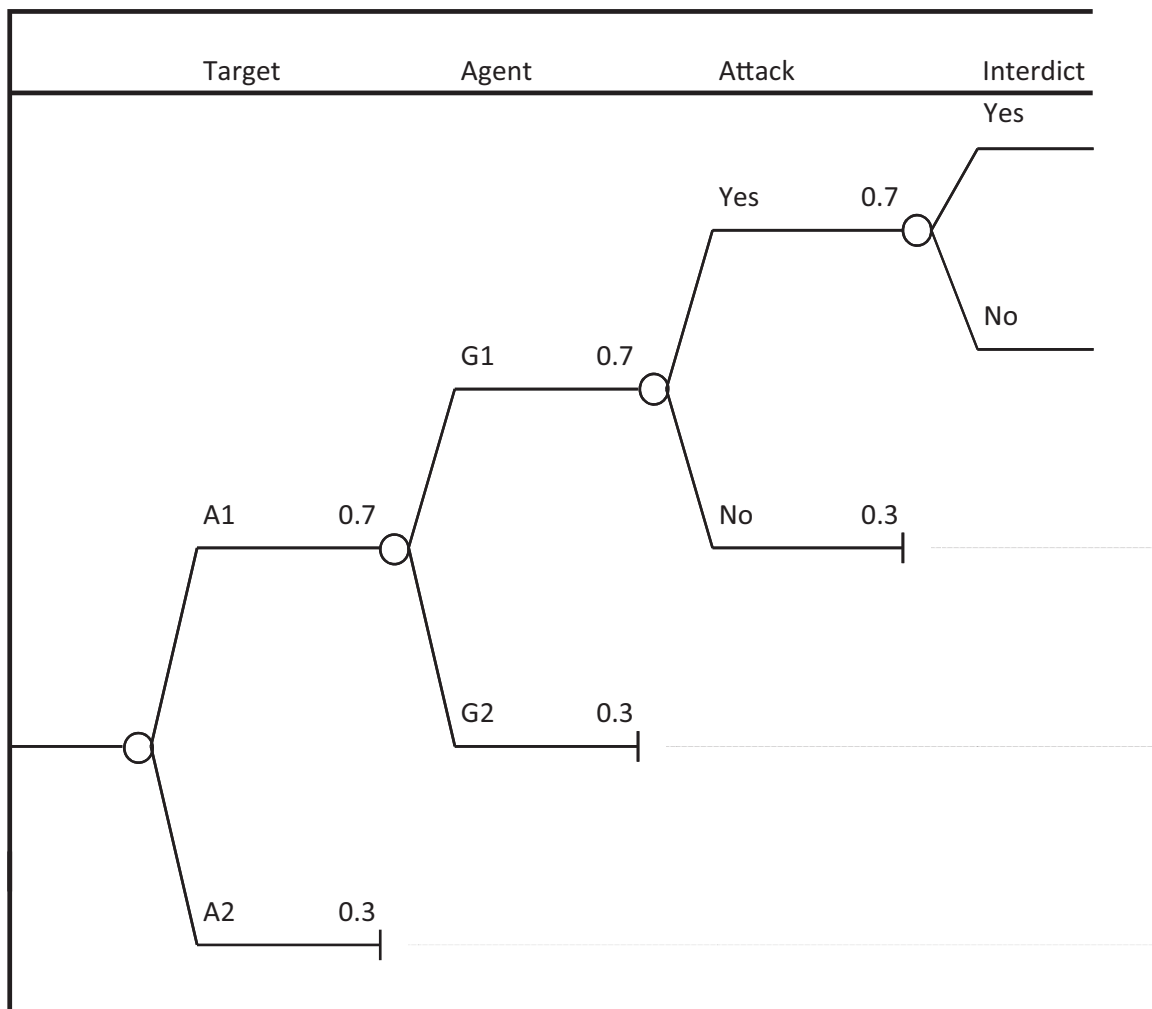


Fig. 3. Notional bioterrorism event tree.

intelligence information describing what the adversary is actually *doing* is uncertain and incomplete, data about the specifics of the objective function a terrorist may be trying to maximize—that is, what the terrorist is *thinking* is at least as uncertain and incomplete, and probably more so. Specifically, this assumption requires knowledge of the adversary's preferences and capabilities relating to

- representational and computational capacities⁽²⁷⁾
- the nature of the adversary's objective function (i.e., whether the goal is to maximize his or her consequences);
- adversary knowledge and reasoning about the problem commensurate with the United States; and

- adversary capability of representing a model and computing its solution.⁽²⁸⁾

To summarize, probability trees, event trees, and decision trees share common elements of forward logic and appear very similar. Event trees have a long history in PRA and in the past 10 years many applications to terrorist risk have been published. Applications using decision trees in terrorism risk, however, are not well developed in the risk analysis literature. In comparing Figs. 3 (event tree) and 4 (decision tree from an attacker's perspective), the distinction between the two formulations becomes clear. In Fig. 3, the event tree formulation expresses *the intelligence community's uncertainty in the intelligent adversary's true utility*. This uncertainty is represented by a probability assignment for a terrorist

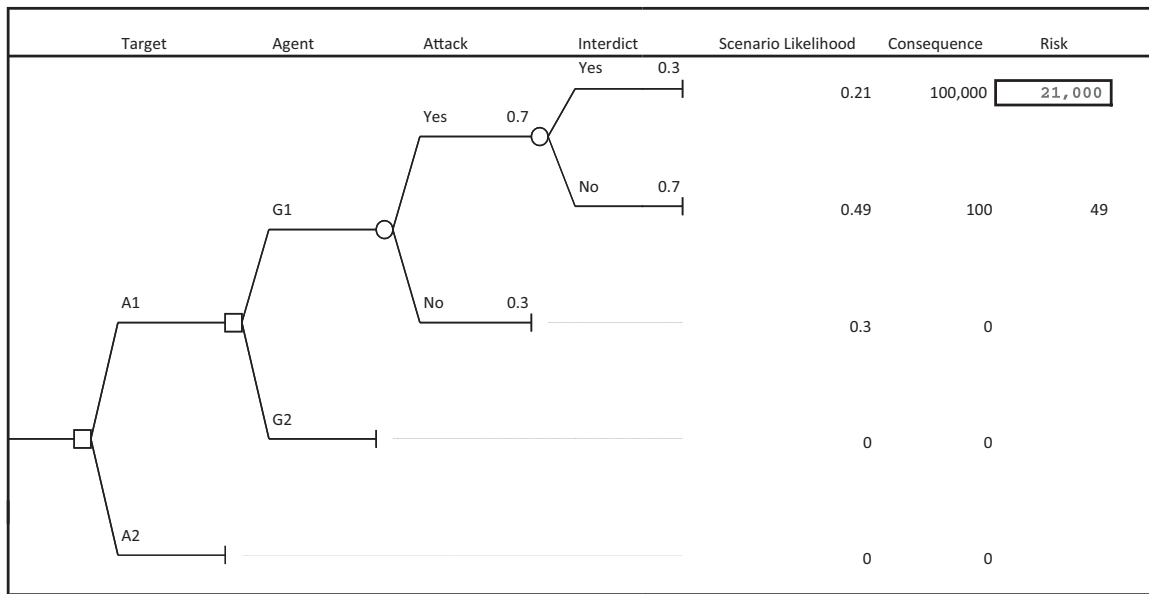


Fig. 4. Notional bioterrorism decision tree.

choice based on the state of information available to the intelligence community. In Fig. 4, and as described above, the decision tree assumes (1) we know the intelligent adversary's true utility for the choices in the tree, and (2) the adversary has the knowledge and rationality to actually maximize that utility; we argue that this is not the case.

3.1.2. Fault, Attack, and Success Tree

Fault trees deductively model the combinations of system failures and human errors that could lead to an accident (Fig. 5). Fault trees are a useful tool for analyzing, visually displaying, and evaluating failure paths in a system to evaluate system level risk.⁽²⁹⁾ Similarly, success trees model the combinations of events that lead to success. Attack trees model the actions of an intelligent adversary to defeat a defensive system.⁽³⁰⁾ Fault, attack, and success trees all use Boolean algebra, reliability theory, and probability theory. Fault trees provide insight into why mitigating events in an event tree may fail. Attack trees categorize the different ways in which a system can be attacked. And, success (defense) trees are the complement of fault trees. Fault trees have been used for years in concert with event trees. For example: 1961—Minute Man Missile, aircraft design, and nuclear power plant safety analyses have all used fault trees to support planning and decisions. Similarly, attack trees are special case of fault trees that have been used to represent an adversary's successful de-

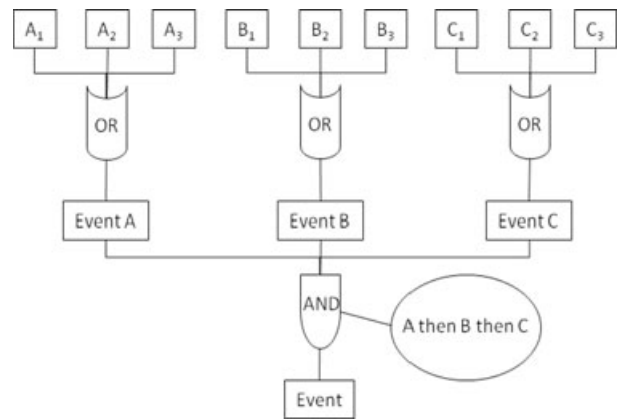


Fig. 5. Generic fault tree.

feat of a countermeasure or system, such as a firewall system, a mitigating event for an event tree.⁽³¹⁾ Along this same line, other techniques may prove to be useful with logic trees. For example, in the next section, a combination of systems dynamic modeling and Bayesian networks is explored to provide more structure to how one may underpin estimates from the intelligence community for the behavior of the intelligent adversary.

3.2. Influence Diagrams

Influence diagrams depict the relationships between decisions, events, and outcomes by

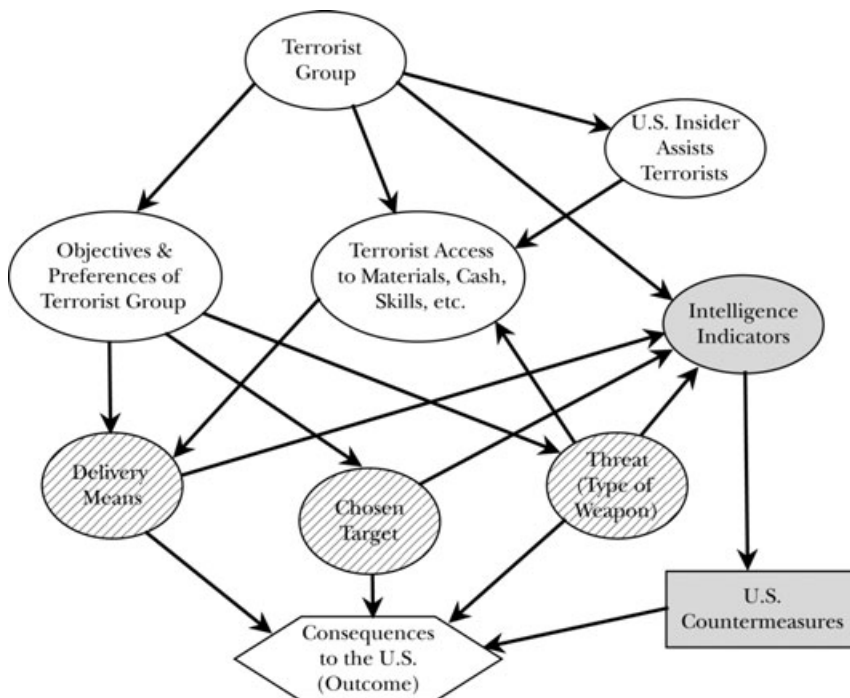


Fig. 6. Example terrorist influence diagram.

depicting them as nodes and arcs in a directed acyclic graph (Fig. 6). Following influence diagram conventions, uncertain variables and events are typically shown as ellipses, variables calculated from predecessors are shown as double ellipses, decision nodes are shown as rectangles, and value nodes are shown as diamonds (or hexagons). Influence diagrams are directed acyclic graphs, in which an arrow connecting node A to node B is interpreted as follows:

1. If node A is an event node and node B is a decision node, it means that “the event in A will be known prior to making the decision in D.”
2. If both nodes A and B are both chance nodes, it means that “knowing the event in node A affects the probabilities of events in node B.”
3. If node A is a chance or decision node and node B is an outcome node, it means that “the outcome depends on the predecessor nodes.”
4. If nodes A and A' are chance nodes and node B is a node characterizing a calculated variable, it means that the variable is node B is calculated from the numbers representing the uncertain variable in A and A'.

Influence diagrams are also compact representations of decision trees, which hide the branches of the

tree and display only its nodes. Every influence diagram can be represented by a symmetric decision tree, in which all possible combinations of decision, events, and outcomes are represented. Conversely, all decision trees can be enriched to symmetric trees and converted to an influence diagram.

Fig. 6 shows an influence diagrams for a terrorist group’s decision to choose a mode and a target of attack.⁽³²⁾ In this influence diagram, there is only one decision node, characterizing the possible countermeasures, and one outcome node, the consequences of a combination of the type of weapon, target, and delivery means (shaded event nodes in Fig. 6). The other nodes are event nodes, characterizing terrorist states and information available to the U.S. government.

This influence diagram analysis was applied in an illustrative analysis to determine the relative likelihood of alternative types of terrorist attacks on the United States.

The results suggested that IED attacks, “dirty bombs,” and biological attacks are most likely.

3.3. Causal Loop Diagrams and Systems Dynamic Models

Representing terrorist behavior in a simulation is a unique challenge. When simulating physical

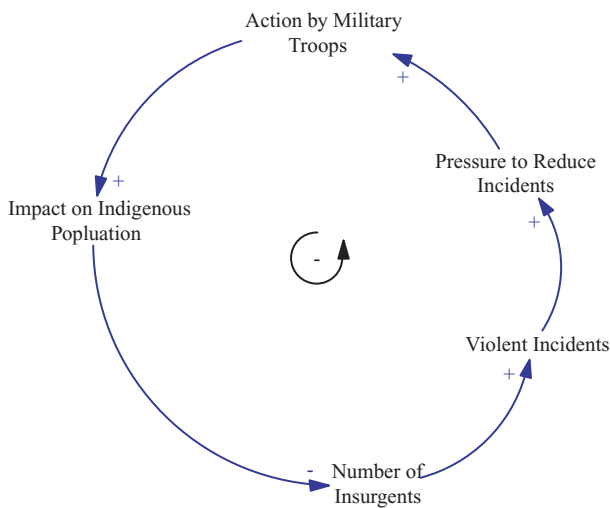


Fig. 7. Insurgency causal loop diagram.

phenomena one can utilize detailed mathematical models to capture the behavior of those systems. Models such as those representing flight dynamics or fluid structures are well understood and have been developed to a point where they closely represent reality. Some representations have become good enough to substitute the models and simulations for real-world systems and still attain the same or nearly the same training and testing benefits as the actual system. The same cannot be said for the terrorist and social systems. Human systems are far more complex than physical systems and therefore much more difficult to represent in computational models. However, strides have been made in this area and a few points are worth mentioning.

A modeling paradigm that lends itself to capturing the many factors that influence complex system behavior is system dynamics (SD). SD is a way to model and depict factors contributing to the behavior of a system and the causal relationships that exist among those factors. To employ this modeling method one begins by developing a causal loop diagram that shows these variables and their causal linkages. Fig. 7 represents one such causal loop diagram.

This figure is a simple example of factors that are causing a change in an insurgent population. The words in the diagram represent the factors linked to the insurgent level and the arrows show the causal linkages. A plus sign indicates a change in the same direction as the source factor change. A minus sign indicates a change in the opposite direction. One can interpret this diagram as follows. As the num-

ber of insurgents change, this causes a change in the number of violent incidents. As those change, pressure on the government to reduce incidents changes. This change causes a subsequent response by military troops, which leads to a change in how the indigenous population responds. Their response then contributes to the number of people participating in the insurgency.

Once the causal loop diagram is developed the modeler can proceed to the next step in SD, which is the development of a stock and flow diagram. This type of diagram is derived from the causal loop relationships and represents actual levels and rates of change of the system variables. Fig. 8 is the stock and flow representation for this insurgent model.

The rectangular box is the level variable of interest that is governed by creation and loss rates. These rates are in turn influenced by the factors noted in the causal loop diagram. With an SD model one has a graphical representation of the system that clearly shows the cause and effect relationships and provides a sense of the interdependencies that exist in complex systems.

Although a simple example, one can see that this approach lends itself to modeling many types of systems. By including variables that affect the risk level of a particular problem, a sense of how changes in those variables will affect the risk level may be gleaned.

3.4. Bayesian Network Analysis

A Bayesian network is a directed acyclic graph or belief network where nodes represent random variables and the directed arcs indicate probabilistic dependence. The arcs in the network define probabilistic dependence between pairs of variables; the direction of the arc indicates which of two possible conditional probability distributions has been captured.⁽³³⁾ In describing a network, modelers use tree and family metaphors to describe the relationships between nodes. Parents are nodes that point to children. Ancestors of a node include all nodes that point to it directly or indirectly through its parents. Descendants are those nodes to which a node points directly or indirectly.⁽³⁴⁾

Bayesian networks have been used in the development of anti-terrorism modeling. BN have also been used to predict distribution for lethal exposure to chemical nerve agents such as Sarin. For example, in Fig. 9, conditional probability tables are shown for each node: Sarin Attack, Exposure Type, Detected

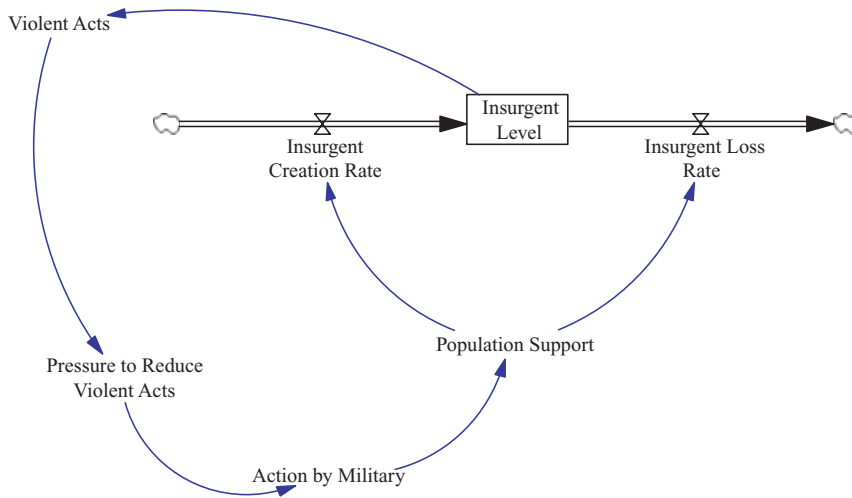


Fig. 8. Insurgent stock and flow diagram.

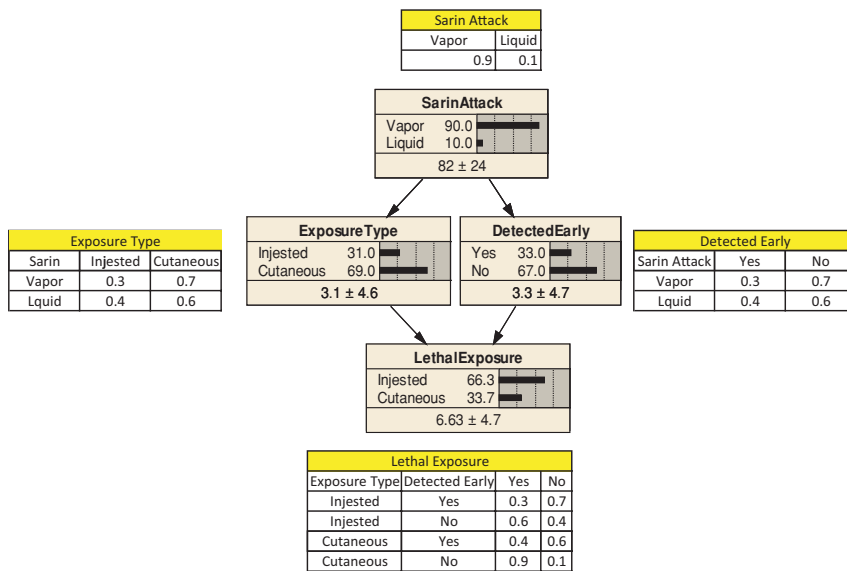


Fig. 9. Notional Sarin nerve agent Bayesian network.

Early, and Lethal Exposure. In this notional model, Sarin Attack is the root node and parent to Exposure Type and Detected Early. By changing the evidence for a given variable, the inference algorithm applies Bayes's rule and updates the probability distribution for Lethal Exposure. BN are useful as well because as the situation changes, they are easy to update; as the evidence changes, the posterior probability changes. BN software such as Netica™ or GeNie™ is well known and can interact and exchange data with other models.

3.5. Game Theoretic Models

Game theory (GT) is the study of multi-agent decision problems. The classic references include von

Neumann and Morgenstern, and Nash, though it was first considered in the 19th century by Cournot.^(35–37) Most of the current research and applications are conducted by micro-economists, though the subject is not confined to that area as there have also been several successful applications of the technique in areas as diverse as computer science and evolutionary biology.^(38,39)

An important assumption in game-theoretic models is that all the possible utilities of the different consequences for each player must be derivable and usable within the model. This implies knowledge of the possible goals and aspirations of the different players. Classical game theory only considers one set of utilities for each player; however, if there is uncertainty about the other players' intentions,

multiple utilities can be modeled using games of incomplete information.^(40–42) In addition, another well-discussed assumption of game theory is the need for the players to be *rational* and *intelligent* enough to work out the consequence of their actions.⁽⁴³⁾ This assumption can be quite limiting especially when players turn out to be less sophisticated than first thought (e.g., they have misinterpreted the consequences of their actions). It has also been questioned whether any sensible decisions can be made about terrorist behavior for analytical application.⁽⁴⁴⁾ However, these limiting assumptions have not prevented game theory from contributing to analysis in numerous related areas.⁽⁴⁵⁾

One response that is often offered in defense of the assumption above, that adversary objectives are strictly consequence maximizers, is that it defaults to the worst case and is therefore “erring” in the right direction. The thinking is that planning or defending against the worst an intelligent agent can do is not necessarily a bad thing. However, there are some limitations to this argument. For example, if resources have to be removed from defending against a small-scale attack to help defend against a large-scale attack this could make attacking on the small scale beneficial to the terrorist *especially* if that was the intention anyway. Further, the worst-case consequences for terrorism scenarios (particularly in the case of WMD attacks) can be many orders of magnitude larger than the assessed expected consequences given less-than-optimal choices made by the adversary. This can greatly bias assessment conclusions and have strong impacts on planning and defense strategies.

There have been several successful applications of game theory within the nuclear, biological, and chemical (NBC) arena. For example, Robert Aumann, Nobel Memorial Prize in Economics 1995, and Michael Maschler developed game theoretic models during the Cold War for dealing with nuclear disarmament problem.⁽⁴⁶⁾ Most of the work using game theory within an NBC arena focuses on the human element of the problem, to include negotiating with terrorists, formation of terrorists groups, and willingness of the domestic population to accept solutions to a terrorism problem.^(47–50) More recent applications of game theory can be seen in Kott and McEneaney, who discuss several applications of the game theory to Department of Defense problems and cover a wide variety of possible implementations.⁽⁵¹⁾

It is important to realize that unlike other analytical methods (where finding the maximum or min-

imum of some value is the sole goal), there is no “one size fits all” solution method to a game. For example, possible solution methods include min-max.^(35,36) Also, these solution methods can produce multiple solutions to the same problem (due to the nonlinear nature of games). This has led to much research into the selection of the solution when multiple solutions are presented.^(52,53)

Recent useful applications of game theory have been used in analyses supporting the protection of multiple targets and to interdependent security.^(54,55) A very important application of game theory was developed by Milind Tambe and his colleagues, who used the Stackelberg game to develop optimal randomizations of inspections and patrols.⁽⁵⁶⁾ Applications are spreading rapidly to include randomizations of patrols and inspections at the Los Angeles Airport to randomization of assignments of Federal Air Marshalls, among other areas.

Game theory is a normative technique as opposed to a descriptive or positive one.⁽⁴³⁾ This means that given a game, game theory will tell you how the game should be played as opposed how it will actually be played. This may lead an analyst to gain some unexpected and interesting insight into the terrorism problem, which other techniques fail to provide. However, given the problems of determining an opponent’s rationale and intentions, a good solution might given by game theory but it could be for the wrong game in the first place.

4. CONCLUSION

There remains significant debate in the risk analysis, decision analysis, security, and intelligence communities about the validity of the NRC Committee’s underlying assumption that terrorists have the understanding to always make optimal choices that maximize consequences in the wide array of challenging technical disciplines required for successful execution of a bioterrorism attack. Despite decades of advances in consequence modeling in the U.S. government, there may be significant uncertainty and inaccuracy in our own estimates of consequences from a WMD attack given a set of initial conditions—Can we base our strategic biodefense planning on the assumption that terrorists do better? That they have the information at their disposal to always make consequence-maximizing choices? Given this, eliciting intelligence experts that have a level of understanding of what our adversaries are capable of accomplishing is a quite reasonable and responsible

option; representing this information probabilistically on an event tree is reasonable as well, noting the relevant limitations.

Professor Yacov Haimes suggests the use of multiple techniques for assessing terrorist actions as probabilities. He reminds us that “no single model or methodology can effectively meet all the challenges of tracking terrorism through scenario generation and structuring, updating and quantifying the value of intelligence, assigning priorities to the scenarios in a well-established risk-based methodology, or track terrorists’ attack plans.”⁽²⁶⁾ We agree that multiple approaches, perhaps in combination are needed to address the complex issue of terrorism, including event trees, decision trees, fault trees, Bayesian belief networks, game theory, and agent-based models, among others. In this context, however, PRA and event trees have been shown to be useful approaches for assessing terrorism risks, especially for creating a baseline comparison of these risks. Decision trees, like PRA and all approaches, have limitations and are not on their own a complete solution. In the case of applications for terrorism risk analysis, the NRC’s decision tree approach has limitations that may be difficult to surmount upon implementation to include the fact that adversaries’ objective functions and level of ability to predict tree outcomes are unknown and difficult to estimate.

ACKNOWLEDGMENTS

The authors would like to thank the contributions of the Area Editor Professor Yacov Haimes and the reviewers for this journal article. Dr. Detlof von Winterfeldt’s participation in writing this article was supported by Decision Research, Oregon and the U.S. Department of Homeland Security under Grant 2008-GA-T8-K004 (Federal Emergency Management Agency, National Preparedness Directorate) to the University of Southern California. Points of view and opinions expressed in this article are solely those of the authors and do not necessarily represent the official position or policies of Decision Research or FEMA’s National Preparedness Directorate or the Department of Homeland Security.

REFERENCES

1. Committee on Methodological Improvements to the Department of Homeland Security’s Biological Agent Risk Analysis, National Research Council (NRC). Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change, 2008. Available at: http://books.nap.edu/openbook.php?record_id=12206.
2. U.S. Department of Homeland Security. Response to National Academies of Science, 2008. Available at: <http://www.dhs.gov/xabout/structure/gc.1222807486868.shtm>.
3. Kaplan S, Garrick BJ. On the quantitative definition of risk. *Risk Analysis* 1981; 1(1):11–27.
4. Garrick B, Hall J, Kilger M, McDonald J, O’Toole T, Probst P, Parker E, Rosenthal R, Trivelpiece A, Arsdale L, Zebroski E. Confronting the risks of terrorism: Making the right decisions. *Reliability Engineering and System Safety*, 2004; 86(2):129–176.
5. Garcia M. Vulnerability Assessment of Physical Protection Systems, Sandia National Laboratories. US: Elsevier, 2006.
6. McGill W, Ayyub B, Kaminskiy M. Risk analysis for critical asset protection. *Risk Analysis* 2007; 27(5):1265–1281.
7. Pate-Cornell ME, Guikema SD. Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research*, 2002 Dec; 7(4):5–23.
8. Rosoff H, von Winterfeldt D. A risk and economic analysis of dirty bomb attacks on the ports of Los Angeles and Long Beach. *Risk Analysis*, 2007; 27(3):533–546.
9. Willis H, Morral A, Kelly T, Medby J. Estimating terrorism risk. RAND Center for Terrorism Risk Management Policy, 2003.
10. von Winterfeldt D, O’Sullivan T. Should we protect commercial airplanes against surface to air missile attacks from terrorists? *Decision Anal.* 2006; 3(2):63–75.
11. Wilson R. Combating terrorism: An event tree approach. Proceedings of the International Seminar on Nuclear War and Planetary Emergencies 27th Session, 2003; 122–143.
12. Bedford T, Cooke, RM. Probabilistic Risk Analysis Foundations and Methods. Cambridge, UK: Cambridge University Press, 2001.
13. Hora S. Eliciting probabilities from experts. In Edwards W, Miles R, Jr., von Winterfeldt D (eds). *Advances in Decision Analysis*. Cambridge, UK: Cambridge University Press, 2007.
14. Apostolakis G. The concept of probability in safety assessments of technological systems. *Science*, 1990 Dec; 250(4986):1359–1364.
15. Kent S. The Board of National Estimates: Collected Essays, Center for the Study of Intelligence, 2007. Available at: www.cia.gov.
16. Ezell B, Haimes Y, Lambert J. Risks of cyber attack to water utility supervisory control and data acquisition systems. *Military Operations Research*, 2001; 6(2):30–46.
17. Koller G. Risk Modeling for Determining Value and Decision Making. FL: Chapman and Hall, 2000.
18. Viscusi WK. *The Risk of Terrorism*. US: Kluwer Academic Publisher, 2003.
19. Dillon-Merrill R, Parnell G, Buckshaw D. Logic trees: Fault, success, attack, event, probability, and decision trees. *Wiley Handbook of Science and Technology for Homeland Security*, 2008.
20. Cornell CA, Merz H. Seismic risk analysis of Boston. *Journal of the Structural Engineering Division*, 1975 Oct; 101(10):2027–2043.
21. Parnell G, Liebe R, Dillon-Merrill R, Buede D, Scouras J, Colletti B, Cummings, M, McGarvey D, Newport R, Vinch P, Ayyub B, Kaminskiy M, Scouras J. *Homeland Security Risk Assessment: Volume 1—An Illustrative Framework and Volume 2—Appendices of Methods*. Washington, DC: Homeland Security Institute, 2005.
22. Kumamoto H, Henley E. *Probabilistic Risk Assessment and Management for Engineers and Scientists*. US: IEEE Press, 1996.
23. Bommer J, Scherbaum F, Bungum H, Cotton F, Sabetta F, Abrahamson N. On the use of logic trees for ground-motion

- prediction equations in seismic-hazard analysis. *Bulletin of the Seismological Society of America*, 2005 Apr; 95(2):377–389.
24. Vesely W. Quantitative risk analysis: Challenges and opportunities at NASA, Manager Risk Assessment for NASA Headquarters, 2005. Available at: ti.arc.nasa.gov/projects/ishem/Papers/Vesely_PRA.doc
 25. Kirkwood C. A Decision Tree Primer, 2002. Available at: www.public.asu.edu/kirkwood/
 26. Goodwin D. *The Military and Negotiation: The Role of the Soldier-Diplomat*. Taylor & Francis, 2005.
 26. Haimes Y. *Risk Modeling, Assessment, and Management*. NY: John Wiley and Sons, 2009.
 27. Hees M, Roy O. *Intentions, Decisions and Rationality*, 2007. Available at: www.ilic.uva.nl/Publications/ResearchReports/PP-2007-21.text.pdf
 28. Cordesman A. *The Challenge of Bioterrorism*. MA: CSIS Press, 2005.
 28. Myerson R. *Game Theory: Analysis of Conflict*. Washington, DC: Harvard UP, 1991.
 29. Ericson C. Fault tree analysis: A history. *Proceedings of the 17th International System Safety Conference*, 1999 Aug.
 30. Mauw S, Oostdijk M. *Foundations of Attack Trees*, 2006; Available at: <http://www.cs.ru.nl/~martijno/publications/papers/attacktrees.pdf>.
 31. Ezell B. Risks of cyber attack to supervisory control and data acquisition for water supply, Thesis, University of Virginia, May 1998.
 32. Pate-Cornell ME. The engineering risk analysis method and some applications. In Edwards W, Miles R, Jr., von Winterfeldt D (eds). *Advances in Decision Analysis*. Cambridge, UK: Cambridge University Press, 2007.
 33. Spirtes P, Glymour C, Scheines R, Kauffman S, Aimalé V, Wimberly F. Constructing Bayesian network models of gene expression networks from microarray data. *Proceedings of the Atlantic Symposium on Computational Biology*, 2000.
 34. Hudson L, Ware B, Laskey K, Mahoney S. An application of Bayesian networks to antiterrorism risk management for military planners. C4ISR Papers Mason Archival Repository, 2005. Available at: hdl.handle.net/1920/268.
 35. von Neumann J, Morgenstern O. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1944.
 36. Nash J. Non-cooperative games. *Annals of Mathematics*, 1951; 54(2): 286–295.
 37. Cournot A. *Recherches sur les principes mathématiques de la thorie des richesses (Researches into the mathematical principles of the theory of wealth)*. Hachette, 1838.
 38. Dash RK, Jennings NR, Parkes DC. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 2003 Dec; 18(6):40–47.
 39. Smith M. The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 2004; 47(1):209–221.
 40. Harsanyi J. Games with incomplete information played by Bayesian players I. *Management Science*, 1967; 14(20):159–182.
 41. Harsanyi J. Games with incomplete information played by Bayesian players II. *Management Science*, 1968a; 14(20):320–334.
 42. Harsanyi J. Games with incomplete information played by Bayesian players III. *Management Science*, 1968b; 14(20):486–502.
 43. Binmore K. *Essays on the Foundations of Game Theory*. Basil Blackwell, 1990.
 44. Plaut S. Misplaced applications of economic theory to the Middle East. *Public Choice*, 2004; 118(1):11–24.
 45. Borm P, Hamers H, Hendrick R. Operations research games: A survey. *TOP*, 2001 Dec; 9(2):139–199.
 46. Aumann RJ, Maschler MB. *Repeated Games with Incomplete Information*. MA: MIT Press, 1995.
 47. Lapan H, Sandler, T. To bargain or not to bargain? That is the question. *American Economic Review*, 1988; 78(2):16–21.
 48. Anderton CH, Carter JR. Applying intermediate microeconomics to terrorism. *Journal of Economic Education*, 2006 Fall; 37(4):442–458.
 49. Douglas M, Mars G. Terrorism: A positive feedback game. *Human Relations*, 2003; 56(7):763–786.
 50. Bauch CT, Galvani AP, Earn DJ. Group interest versus self-interest in smallpox vaccination policy. *Proceedings of the National Academy of Sciences of the USA*, 2003 Aug; 100(18):10564–10567.
 51. Kott A, McEneaney WM (eds). *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*. FL: Chapman and Hall, 2006.
 52. Harsanyi J, Selten R. *A General Theory of Equilibrium Selection in Games*. MA: MIT Press, 1998.
 53. Herings PJ, Mauleon JA, Vannetelbosch VJ. Fuzzy play, matching devices and coordination failures. *International Journal of Game Theory*, 2004; 32(4):519–513.
 54. Hausken K, Bier VM, Zhuang J. *Defending Against Terrorism, Natural Disaster, and All Hazards, Game Theoretic Risk Analysis of Security Threats*. Springer, 2009.
 55. Heal G, Kunreuther H. You only die once: Interdependent security in an uncertain world. In Richardson HW, Gordon P, Moore JE II (eds). *The Economic Impacts of Terrorist Attacks*. Cheltenham, UK: Edward Elgar, 2005.
 56. Jain M, Pita J, Tambe M, Ordóñez F, Paruchuri P, Kraus S. *Bayesian Stackelberg Games and Their Application for Security at Los Angeles International Airport*, 2008. Available at: <http://www.sigecom.org/exchanges/volume.7/2/jain.pdf>.