2

Probabilistic Risk Analysis and Terrorism Risk

Barry Ezell, Steven P. Bennett, Detlof von Winterfeldt, John Sokolowski, and Andrew J. Collins

2.1 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 Homeland Security's (DHS) Biological Agent Risk Analysis (referred to hereafter as the NRC

Reprinted by permission, Ezell, B.C., Bennett, S.P., von Winterfeldt, D., Sokolowski, J. & Collins, A.J. (2010), Probabilistic risk analysis and terrorism risk. *Risk Analysis*, 30, pp. 575–589.

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

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 chapter is organized in the following way. Section 2.1 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 2.2 details the usefulness of probabilities in bioterrorism risk analysis. Section 2.3 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 2.3 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 decision makers.

2.1.1 *2006 Bioterrorism Risk Assessment Background*

Signed in 2004, Homeland Security Presidential Directive[1] (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.[2] 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" (NRC, 2008).

DHS identified several concerns with the NRC Committee's report (U.S. Department of Homeland Security, 2008). 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 the risk analysis community. This article challenges the NRC Committee's conclusion.

2.1.2 *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 (NRC, 2008).

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.

2.2 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 (Kaplan & Garrick, 1981). 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 and colleagues (2004) advocate the use of PRA for assessing terrorism risk, specifically for assessing the probabilities of terrorist attacks. Work based on Garcia, McGill and colleagues, 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 (Garrick et al., 2004; Garcia, 2006; McGill, Ayyub, & Kaminskiy, 2007; Pate-Cornell & Guikema, 2007; Willis, Morral, Kelly & Medby, 2003; von Winterfeldt & O'Sullivan, 2006).

Willis, McGill and colleagues, 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 (2.1), then, is a common expression of homeland security risk (McGill et al., 2007; Wilson, 2003)

$$Risk = P(A) \times P(S \mid A) \times C \tag{2.1}$$

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 (2.1)) 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 (Bedford & Cooke, 2001; Hora, 2007).

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 (1990) 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."

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 (2007) of the National Board of Estimates replied "65 to 35 in favor of an attack." 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 (2006) 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. 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 decision makers alongside, or in addition to, baseline risks rather than instead of them.

2.3 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 (Ezell, Haimes, & Lambert, 2001; Koller, 2000; Viscusi, 2003). 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 tees to quantify the costs and benefits of countermeasures to MANPADS (Rosoff & von Winterfeldt, 2007; von Winterfeldt & O'Sullivan, 2006).

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.

2.3.1 *Logic Trees*

Logic trees are important tools for exploring the scenario space, analyzing uncertain events, defining scenarios, and assessing risk (Dillon-Merrill, Parnell, & Buckshaw, 2008). 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 (Cornell & Merz, 1975). 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 (Parnell et al., 2005). In this chapter, 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 tress) that the NRC Committee strongly recommended as the appropriate way to do bioterrorism risk analysis.

2.3.1.1 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 (Figure 2.1). 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

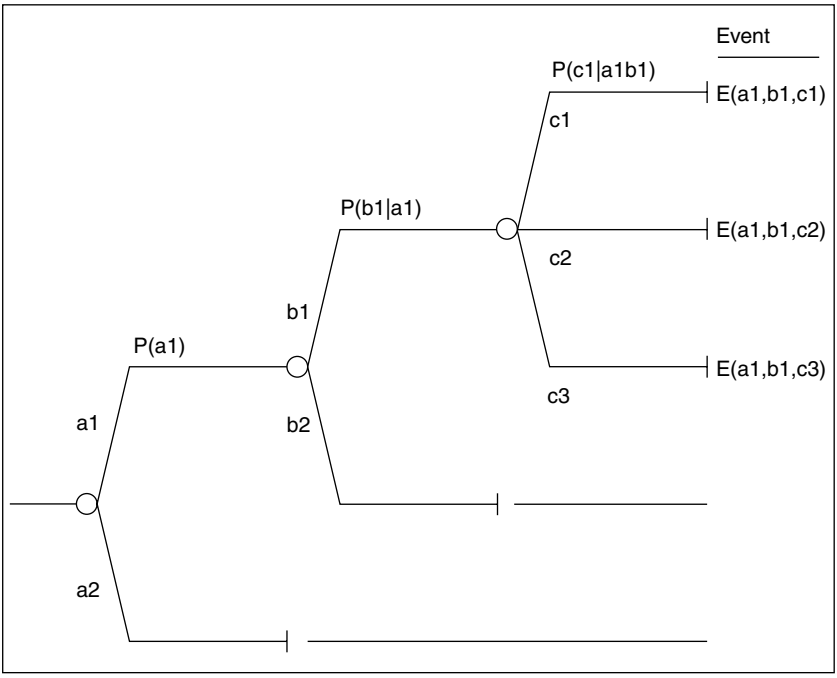

Figure 2.1. Probability tree.

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 (Kumamoto & Henley, 1996). Event trees 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 and colleagues employed event trees to understand cyber risk to supervisory control and data acquisition systems for water supply (Ezell et al., 2011). 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

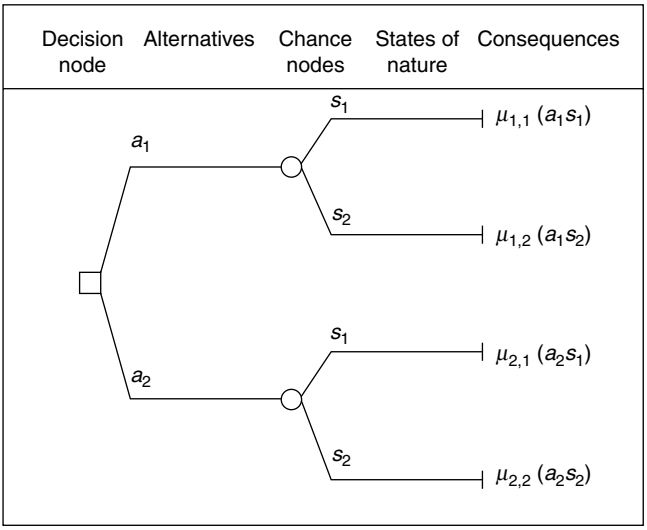

Figure 2.2. Decision tree (Haimes, 2009).

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 (Bommer, Scherbaum, Bungum, Cotton, Sabetta, & Abrahamson, 2005).

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 representation 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 (Vesely, 2005).

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 (Kirkwood, 2002). 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 (Figure 2.2). In the same fashion as probability trees and event