A Comparative Analysis of PRA and Intelligent Adversary Methods for Counterterrorism Risk Management

Greg Parnell
US Military Academy

Jason R. W. Merrick
Virginia Commonwealth University

Simple Counter-terrorism Decision

- Defender has n alternatives
  $$D = \{d_1, d_2, ..., d_n\}$$
- Attacker has m alternatives
  $$A = \{a_1, a_2, ..., a_m\}$$
- States of information
  - Common ε^0
  - Defender ε^D
  - Attacker ε^A
- Consequences
  $$c(d, a), \text{ for } d \in D, a \in A$$

Simple Counter-terrorism Decision

- Usual framework is threat, vulnerability, and consequence
- Vulnerability is represented by two events
  - $P(DS \mid \varepsilon)$ Success or failure of any defenses the defender decides to implement
  - $P(AS \mid \varepsilon)$ Attacker's success in carrying out the attack if the defender fails
- Threat is also represented by two events
  - $P(AC \mid \varepsilon)$ Attacker gains the capability to perform a given attack
  - $A = \{a_1, \ldots, a_m\}$ Attacker decides to use a given attack capability

Container Security

Nuclear materials and devices

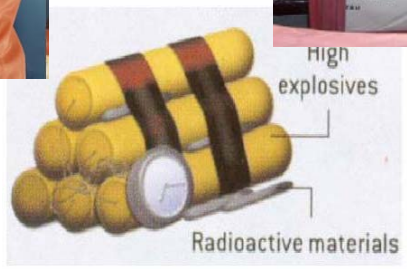

Example 1

- We use container screening for radiological material as a rich and representative example of homeland security decisions
  - Each year 20.4 million containers enter the US (Bonner 2005)
  - Bakir (2008) estimates the probability that terrorists will use a container to smuggle radiological material into the US in the next 10 years to be 0.1. This estimate includes
    - A probability of 0.4 that terrorists will acquire the capability for a RDD
    - A probability of 0.25 that they will attempt to smuggle their device into the US inside a container
  - The probability that the smuggling attempt is thwarted by screening is 0.8
  - Bakir (2008) assumes a 0.5 chance that the attack either is stopped inside the country or is not successfully carried out
  - The consequence distribution for attack 1 to be
    - $10 billion with probability 0.3
    - $40 billion with probability 0.4
    - $100 billion with probability 0.3
  - For attack 2, we assume the consequences to be half these estimates.

Defender Event Tree

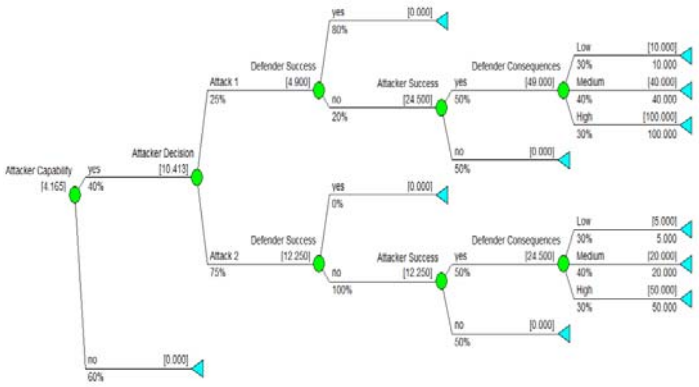

$$P(AC \mid \varepsilon^D)\sum_{j=1}^{2}\Big(P(A=\mathrm{a}_j \mid \varepsilon^D)P(\overline{DS} \mid D=d_1, A=\mathrm{a}_j, \varepsilon^D)P(AS \mid D=d_1, \varepsilon^D)E[c(d_1,a_j) \mid \varepsilon^D]\Big).$$

Defender Bayesian Network

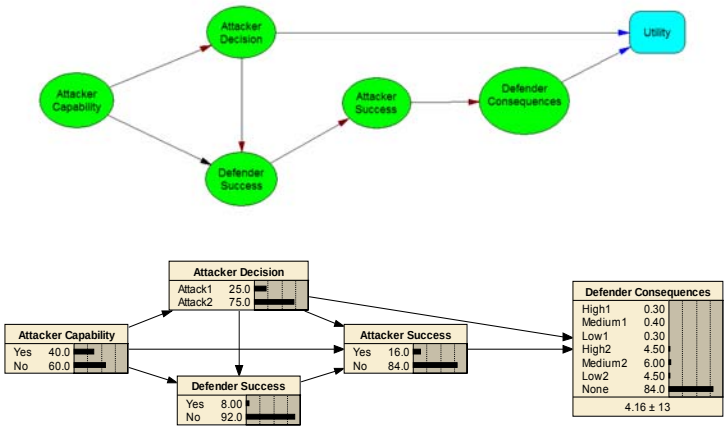

Attacker Event Tree

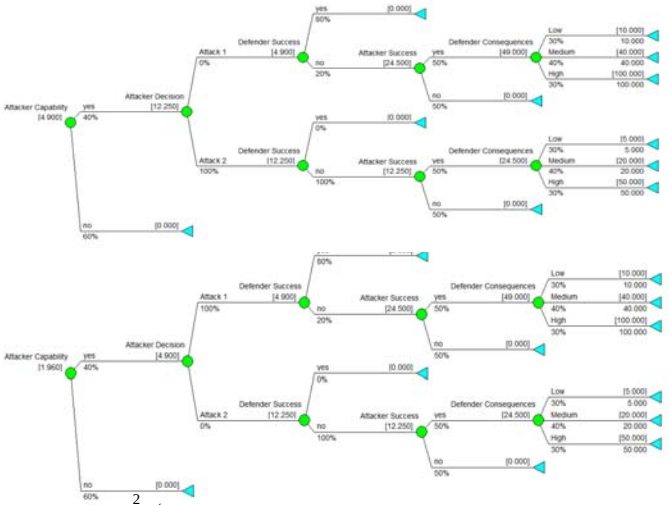

Attack 2

Attack 1

$$P(AC \mid \varepsilon^D)\sum_{i=1}^{2}\left(P(\overline{DS} \mid D=d_i, A=\mathrm{a}_j, \varepsilon^D)P(AS \mid D=d_i, \varepsilon^D)E[c(d_i,a_j) \mid \varepsilon^D]\right)$$

Attacker Bayesian Network

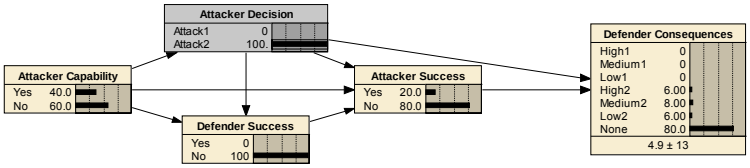

Attack 2

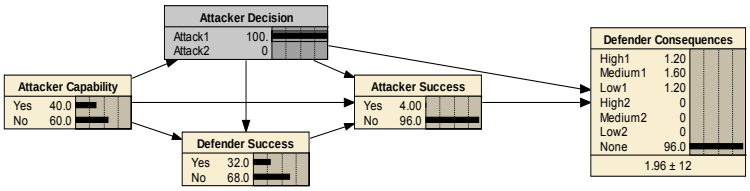

Attack 1

Defender Decision Tree

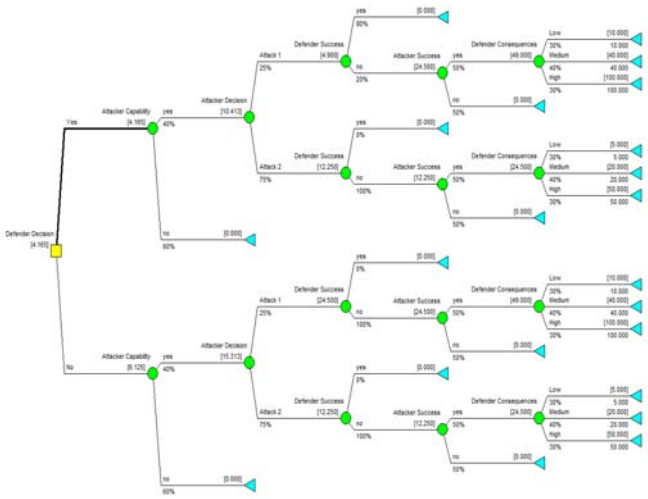

$$d^* = \arg\min_{d \in D} \left\{ P(AC \mid \varepsilon^D) \sum_{j=1}^{2} \left(P(A = \mathrm{a}_j \mid \varepsilon^D) P(\overline{DS} \mid D = d, A = \mathrm{a}_j, \varepsilon^D) P(AS \mid D = d, \varepsilon^D) E[c(d, a_j) \mid \varepsilon^D] \right) \right\}$$

Attacker Decision Tree

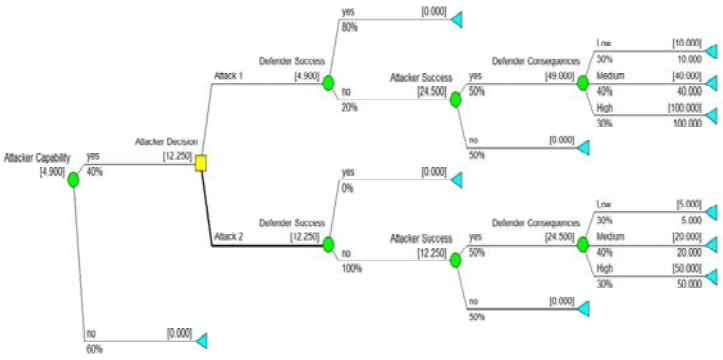

$$a^* = \arg\min_{a \in A} \left\{ P(AC \mid \varepsilon^D) P(\overline{DS} \mid D = d_1, A = \mathrm{a}, \varepsilon^D) P(AS \mid D = d_1, \varepsilon^D) E[c(d_1, a_j) \mid \varepsilon^D] \right\}$$

Simultaneous Games

		Attacker Decision: Attack 1	Attacker Decision: Attack 2
Defender Decision	Yes	1.96	4.9
	No	9.8	4.9

$$d^* = \arg\min_{d \in D} \left\{ \bar{c}(d, a^*) \right\} \qquad a^* = \arg\min_{a \in A} \left\{ \bar{c}(d^*, a) \right\}$$

Sequential Games

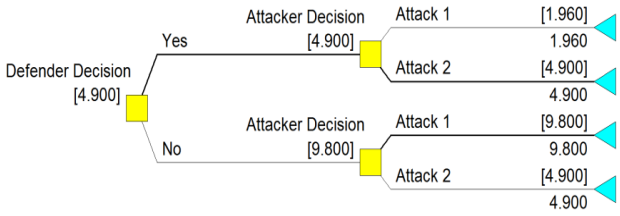

$$a^*(d) = \arg\min_{a \in A} \left\{ \bar{c}(d, a) \right\} \qquad d^* = \arg\min_{d \in D} \left\{ \bar{c}(d, a^*) \right\}$$

Intelligent Adversary Risk Analysis

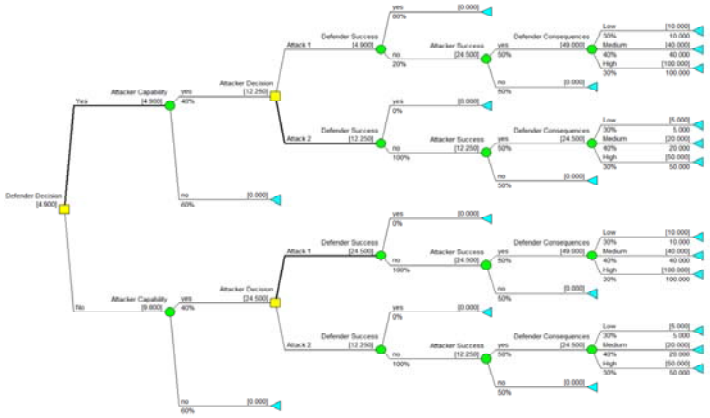

$$a^*(d) = \arg\min_{a \in A} \left\{ P(\overline{DS} \mid D = d, A = \mathrm{a}, \varepsilon^D) P(AS \mid D = d, \varepsilon^D) E[c(d,a) \mid \varepsilon^D] \right\}$$

$$d^* = \arg\min_{d \in D} \left\{ P(AC \mid \varepsilon^D) P(\overline{DS} \mid D = d, A = \mathrm{a}^*, \varepsilon^D) P(AS \mid D = d, \varepsilon^D) E[c(d,\mathrm{a}^*) \mid \varepsilon^D] \right\}$$

Joint Attacker and Defender Decision Trees

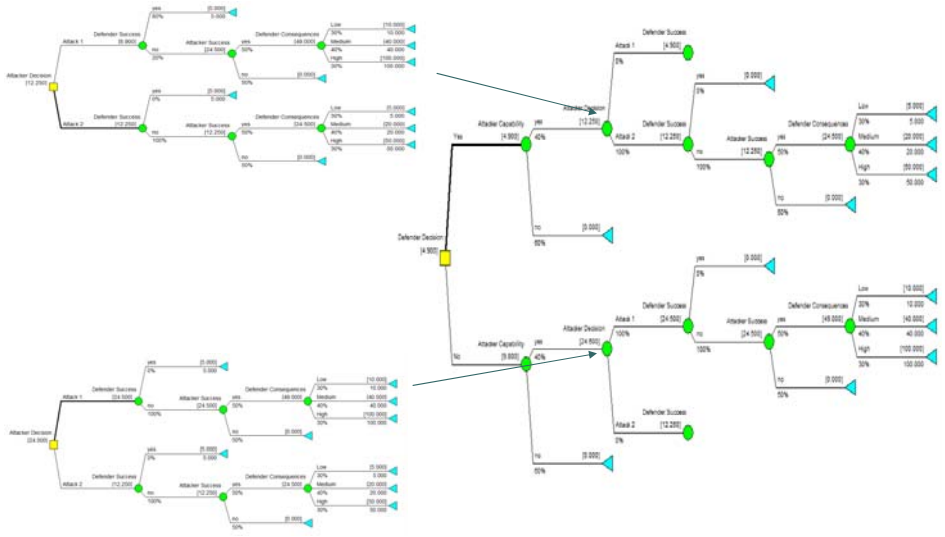

Adversarial Risk Analysis

Attacker sub-model for $D=d_1$

$P(\theta=10) = 0.5$ $\quad$ $P(\theta=5) = 0.5$

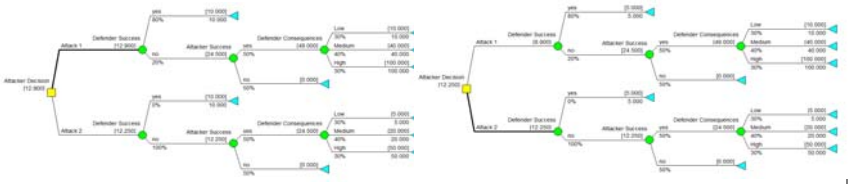

$$a^*(d,\theta) = \arg\min_{a \in A} \left\{ P(\overline{DS} \mid D=d, A=\mathrm{a}, \theta) P(AS \mid D=d, \theta) E[c(d,a) \mid \theta] \right\}$$

$$p^D(a(d) \mid \varepsilon^D) = \int I(a^*(d,\theta) = a(d)) \pi^D(\theta) d\theta$$

Adversarial Risk Analysis

Attacker sub-model for $D=d_2$

$P(\theta=10) = 0.5$ $\quad$ $P(\theta=5) = 0.5$

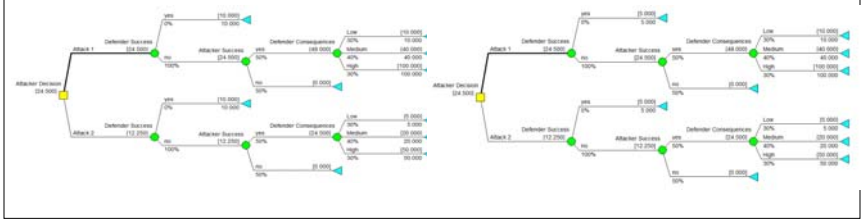

$$a^*(d,\theta) = \arg\min_{a \in A} \left\{ P(\overline{DS} \mid D=d, A=\mathrm{a}, \theta) P(AS \mid D=d, \theta) E[c(d,a) \mid \theta] \right\}$$

$$p^D(a(d) \mid \varepsilon^D) = \int I(a^*(d,\theta) = a(d)) \pi^D(\theta) d\theta$$

Adversarial Risk Analysis

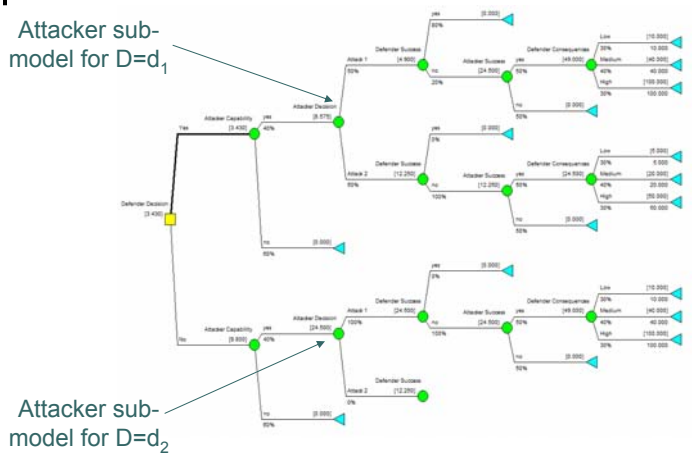

$$d^* = \arg\min_{d \in D} \left\{ P(AC \mid \varepsilon^D) \sum_{a(d) \in A} \left(p^D(a(d) \mid \varepsilon^D) P(\overline{DS} \mid D = d, A = a(d), \varepsilon^D) P(AS \mid D = d, \varepsilon^D) E[c(d, a(d)) \mid \varepsilon^D] \right). \right\}$$

Summary of Methods

Method	Uncertainties	Defender Decisions	Attacker Decisions	State of Information
Defender Event trees	Attacker decision, Attacker capability, Defense success, Attack success given defense failure, Defender consequences	Known a priori	None	Defender's probabilities and consequences used
Attacker Event Tree	Attacker capability, Defender success, Attack success, Attacker consequences	Known a priori	Known a priori	Defender's probabilities and consequences used
Bayesian Network	Any of the above			Defender's probabilities and consequences used
Defender Decision Tree	Attacker decision, Attacker capability, Defense success, Attack success given defense failure, Defender consequences	Solved by backwards induction (minimizing expected defender consequences)	None	Defender's probabilities and consequences used
Attacker Decision Tree	Attacker capability, Screening success, Attack success, Attacker consequences	Known a priori	Solved by backwards induction (maximizing expected attacker consequences)	Defender's probabilities and consequences used
Influence Diagrams	Any of the above			Defender's probabilities and consequences used

Summary of Methods

Method	Uncertainties	Defender Decisions	Attacker Decisions	State of Information
Simultaneous Games	None	Solved by finding Nash equilibrium		Defender's consequences used
Sequential Games	None	Solved by backwards induction (maximizing attacker consequences and minimizing defender consequences)		Defender's consequences used
Intelligent Adversary Risk Analysis	Attacker capability, Defense success, Attack success given defense failure, Defender consequences	Solved by backwards induction (maximizing expected attacker consequences and minimizing expected defender consequences)		Defender's probabilities and consequences used
Adversarial Risk Analysis	Attacker capability, Defense success, Attack success given defense failure, Defender consequences	Solved by backwards induction (maximizing expected attacker consequences and minimizing expected defender consequences)		Defender's probabilities and consequences used in defender tree and defender's beliefs of attacker's state of information.

Results Comparison

Method	Expected Consequences	Defender Decision	Attacker Decisions
Defender Event trees	4.165	Assumed Yes	NA
Attacker Event Tree	1.96 for Attack 1 4.90 for Attack 2	Assumed Yes	NA
Bayes Nets	Equivalent to Event Trees		
Defender Decision Tree	4.165	Yes	Probabilities elicited
Attacker Decision Tree	4.9	Assumed Yes	Attack 2
Influence Diagrams	Equivalent to Decision Trees		
Simultaneous Games	4.9	Yes	Attack 2
Sequential Games	4.9	Yes	Attack 2
Intelligent Adversary Risk Analysis	4.9	Yes	Attack 2
Adversarial Risk Analysis	4.9	Yes	Attack 2
Adversarial Risk Analysis with Uncertainty	3.43	Yes	Probabilities derived from attacker sub-models

Example 2

- The attacker may not obtain attack capability for both attacks.
- We specify a
  - 0.4 chance that they will get the capability for attack 1
  - 0.1 chance they will get the capability for attack 2.
- Thus, there is a
  - 0.04 probability they will get both capabilities
  - 0.36 probability they will get just attack 1 capability
  - 0.06 probability they will get just attack 2 capability
  - 0.54 probability they will get neither capability

Defender Decision Tree

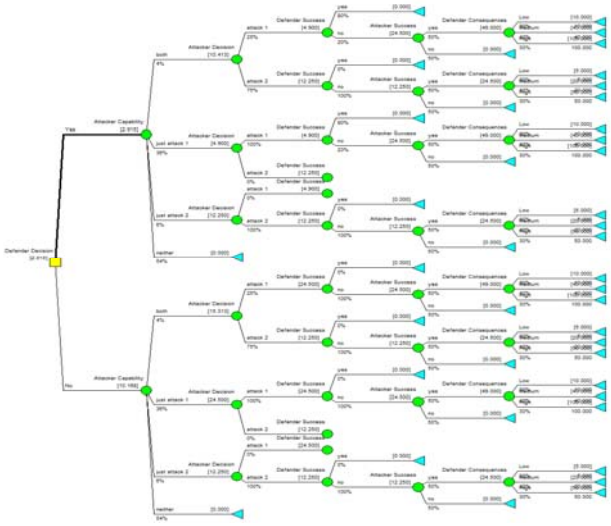

Attacker Decision Tree

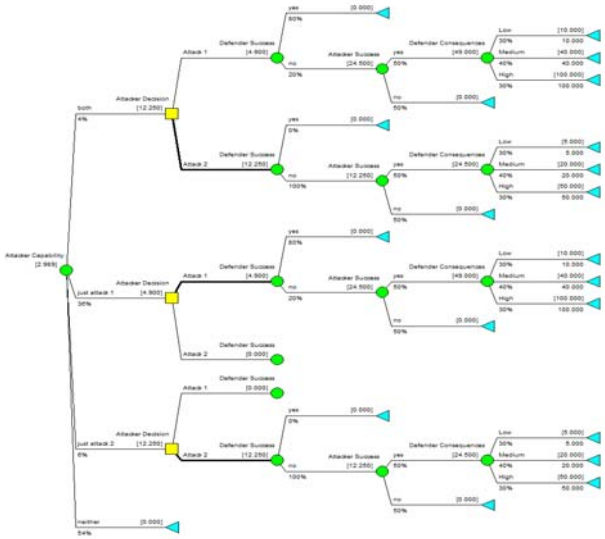

Game Theory

		Attacker Decision: Attack 1	Attacker Decision: Attack 2
Defender Decision	Yes	(1.96)	1.225
	No	9.8	1.225

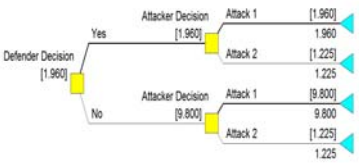

Intelligent Adversary Risk Analysis

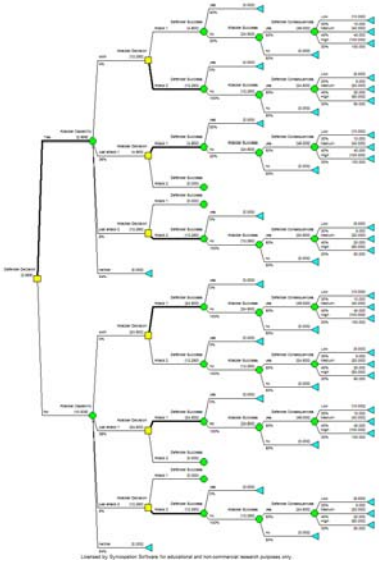

Results Comparison

Method	Expected Consequences	Defender Decisions	Attacker Decisions
Defender Event trees	2.915	Assumed Yes	NA
Attacker Event Tree	2.695 (Attack 1 if capable) 2.989 (Attack 2 if capable)	Assumed Yes	NA
Bayes Nets	Equivalent to Event Trees		
Defender Decision Tree	2.915	Yes	Uncertain
Attacker Decision Tree	2.989	Assumed Yes	Attack 2 if capable
Influence Diagrams	Equivalent to Decision Trees		
Simultaneous Games	1.96	Yes	Attack 1
Sequential Games	1.96	Yes	Attack 1
Intelligent Adversary Risk Analysis	2.989	Yes	Attack 2 if capable
Adversarial Risk Analysis	2.989	Yes	Attack 2 if capable
Adversarial Risk Analysis with Uncertainty	2.842	Yes	Probabilities derived from attacker sub-models

Risk Assessment

- Defender event trees and decision trees that represent attacker decisions as probabilities estimate lower expected consequences than equivalent attacker event trees and decisions trees.
  - Each successive attacker decision modeled as a probability node reduces the expected consequences accordingly.
  - Intelligent adversary risk analysis and adversarial risk analysis estimate expected consequences equal to those of attacker models.
  - Simultaneous and sequential game theory approaches that use simple expected consequence measures do not provide the flexibility to model the order of decisions and uncertainties, and thus arrive at different expected consequences than other approaches in the second example.
- Intelligent adversary risk analysis and adversarial risk analysis estimate the same expected consequences if the same probabilities and consequences are used in the attacker sub-models.

Risk Assessment

- When the defender expresses uncertainty about the attacker's beliefs and preferences in adversarial risk analysis, the result is a distribution over the attacker's decision, the same as specified in the defender decision tree.
  - Ezell et al. (2010) assert that the decision maker must treat the attacker's decision as an uncertainty and specify a probability distribution over the alternatives.
  - However, the adversarial risk analysis approach can be considered a decomposition of this complex distribution, making each elicitation task easier.
  - Suitably decomposed probability elicitations have been found to be better calibrated than those obtained without decomposition (Ravinder et al. 1988, Howard 1989, Mihajlovits and Merrick 2010).

Risk Communication

- Event trees, influence diagrams with just probability nodes, and Bayesian networks with only probability nodes are all equivalent as they are following the laws of probability even though they use different solution algorithms.
  - However, they communicate different aspects of the joint probability distribution.
- Influence diagrams show just the probabilistic dependencies with arrows but do not show probabilities.
- Event trees show prior distributions, conditional probabilities and the expected consequences at each node in the tree.
  - Probabilistic dependencies must be deduced by examining the probabilities in the conditional probability distributions.
  - The marginal probability distributions do not appear in the tree and must be calculated separately.

Risk Communication

- Bayesian networks show probabilistic dependencies (with arrows), the prior distributions and marginal probabilities.
  - They are much more effective for communication since the show all of the prior and marginal probability distributions of the risk results and are easier update as new information becomes available.
- For a risk communications perspective, Bayesian networks seem to be the most useful technique.

Risk Management

- Event trees are less useful for assessing the risk posted by intelligent, adaptive adversaries.
  - Event trees and decisions trees that specify fixed probabilities of attacker decisions will underestimate the potential defender risks.
  - In addition, they do not show the shift in risk given the potential defender decisions.
  - This can also be thought of as the adaptation of the attacker to defender decisions.
- Intelligent adversary risk analysis and adversarial risk analysis explicitly show such adaptation and the shift in risk based on defender decisions.

Are the Attackers Rational?

- We are using these methods prescriptively to improve the defender's decision making
  - Rational decision axioms
  - Use subjective expected utility
- Terrorists are surprisingly quantitative, but they are making unaided decisions
  - Boundedly rational
  - Descriptive decision models