SAND2011-2300C
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

Risk-Based Cost-Benefit Analysis for Security Assessment Problems

Presented at the International Conference on Vulnerability and Risk Analysis and Management (ICVRAM)

April 11-13, 2011 - University of Maryland Conference Center

By Gregory D. Wyss, Ph.D.
Sandia National Laboratories

Co-Authors: **John P. Hinton, Katherine Dunphy-Guzman, John Clem, John Darby, Consuelo Silva, and Kim Mitchiner**

Contact: ☎ (505) 844-5893 gdwyss@sandia.gov

Sandia National Laboratories is a multi program laboratory operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000. Approved for Unlimited Release as SAND2010-xxxxC

Security Risk Management Recommendations from the National Academy of Sciences

- **Our goal must be *effective security risk management*.**

National Academy of Sciences, 2010, emphasis added

Risk management is the process of identifying, analyzing, assessing, and communicating risk and accepting, avoiding, transferring, or controlling it to an acceptable level at an acceptable cost.

- **Key risk management recommendations include:**
  - Use a risk-informed, not risk *based*, approach to security risk management
    - Informed by PRA *tools*, but not relying on PRA
  - Qualitative risk assessment methods may be suitable
  - Focus on risk management rather than "how much or little risk exists"

A Fundamental Definition of Risk

Scenario	Consequence	Likelihood
S_1	C_1	F_1
S_2	C_2	F_2
S_3	C_3	F_3
S_4	C_4	F_4
S_5	C_5	F_5
S_6	C_6	F_6
…	…	…

This *table* *IS* the risk!

– **Risk can be thought of as answers to 3 questions:**

- ***What can happen?*** ***(scenario)***
- ***How likely is it?*** ***(probability / frequency)***
- ***How bad is it?*** ***(consequence)***

"If [a] table contains all the scenarios we can think of, we can then say that it (*the table*) is the answer to the question and therefore *is the risk*."

Kaplan & Garrick, Risk Analysis 1:1(11) 1981, emphasis added.

Risk for a Scenario:

$$R = \underbrace{P_A \cdot (1 - P_E)}_{\text{How likely is it?}} \cdot \underbrace{C}_{\text{How bad is it?}}$$

Applying the Definition of Risk

Scenario	Consequence	Likelihood
S_1	C_1	F_1
S_2	C_2	F_2
S_3	C_3	F_3
S_4	C_4	F_4
S_5	C_5	F_5
S_6	C_6	F_6
…	…	…

This *table* *IS* the risk!

↑ Likelihood / Consequences ➔	Negligible	Low	Moderate	High	Catastrophic
Routine Event					
Unusual Event					
Expected: Life of Facility					
Unlikely: Life of Facility					
Remotely Possible					

Or…

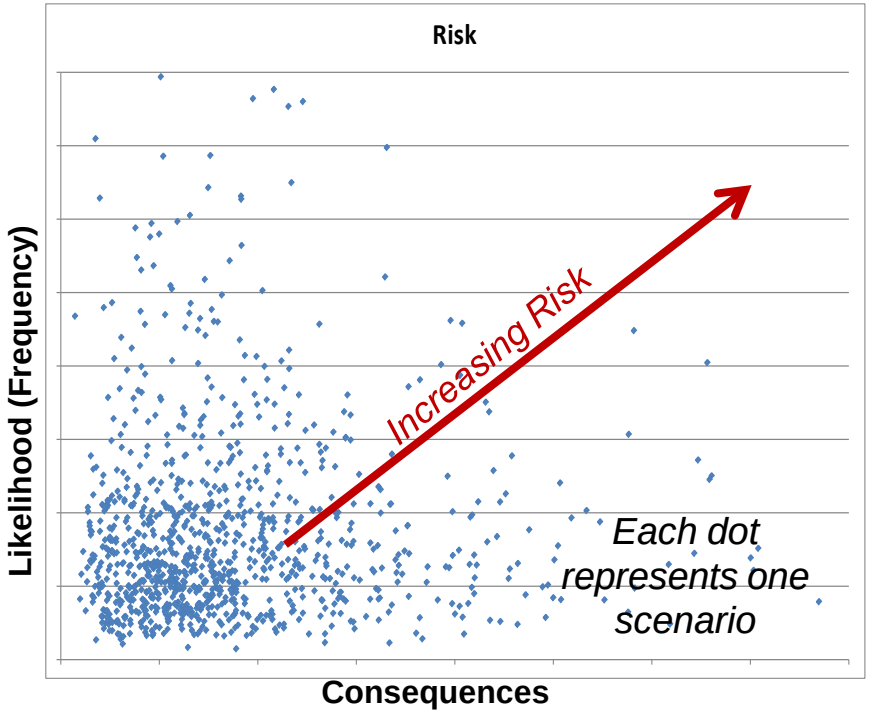

Sandia National Laboratories

The Problem of Likelihood

Attack scenario likelihoods are often elicited from experts.

- **Often assumed by the experts to be statistically independent. *But…***
- **Highly dependent on attacker's capability, motivation & intent**
- **Highly dependent on attacker's other opportunities – both inside and outside the system.**

Security risk estimates are captive to uncertain likelihoods.

- **Which of these is the highest risk?**
- **Which should we mitigate?**
- **Even if we could draw conclusions from this risk picture, the attack likelihood changes frequently and in ways that we may not understand.**

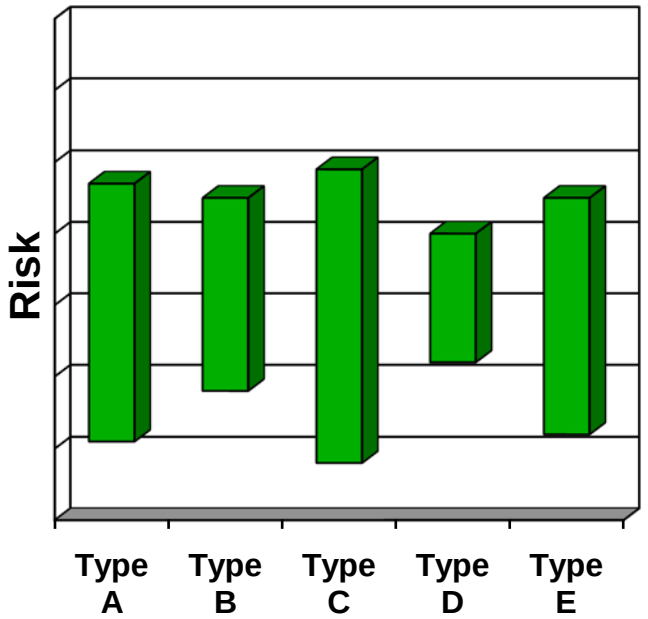

Attack frequency should be the *output* of a risk assessment, not an input.*

** Cox, L.A., Game Theory and Risk Analysis, Risk Analysis, Vol. 29, No. 8, 2009.*

Goal: Manage Security Risks

- **Given uncertainties in attack likelihood, it's hard to get statistically significant recommendations for risk management.**
  - **Can we reduce uncertainty in likelihood?** ***Probably not enough.***

- **A different approach: examine adversary criteria for selecting which attack scenario to pursue, including:**

Adversary's Decision Criterion	How we make an attack less likely
"Could I do it if I wanted to?" ***(Is success likelihood high?)***	**Make attack scenario more difficult**
"Would I do it if I could?" ***(Worthy investment of resources?)*** ***(Does it violate my doctrine?)***	**Make attack scenario more difficult or reduce potential consequences**
"Are the expected consequences high enough?"	**Reduce the potential or expected consequences of the scenario**

- **The benefits of a security investment can be inferred from two metrics:**
  - **How much harder has the scenario become for an adversary?**
  - **How much have expected consequences been reduced?**

Scenario Difficulty Measures the Benefit of a Security Investment

*Illustration based on sites assumed to have the **same consequence** for a successful attack.*

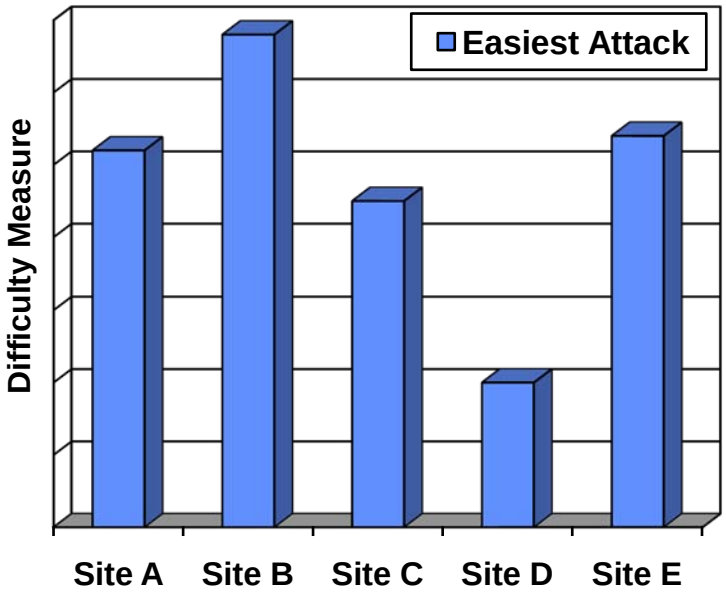

- Are sites balanced?
- Where should I spend my next dollar?

- How much have I improved?
- Why do my sites not meet the new security goal?
- Does this security goal serve the function of a Design Basis Threat?

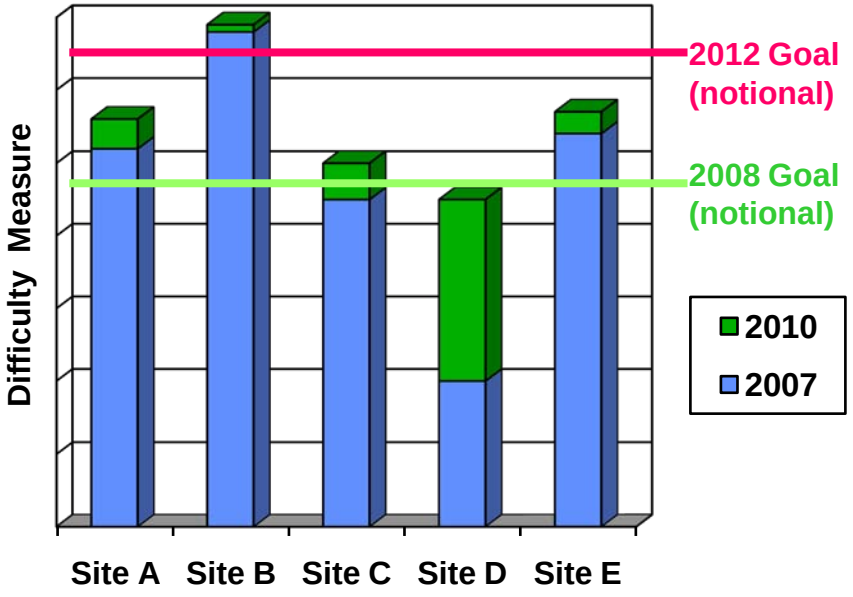

Game theory predicts that, given similar consequences, easier attacks are more likely.

"Scenario difficulty" may be a reasonable surrogate for attack likelihood.

Problems of this type are amenable to traditional optimization methods.

Scenario Difficulty Measures the Benefit of a Security Investment

LDRD
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

If we fix this…

Without fixing this…

We may not have improved security. *Because…*

Many scenarios still exist that are both easier to achieve AND provide higher consequences!

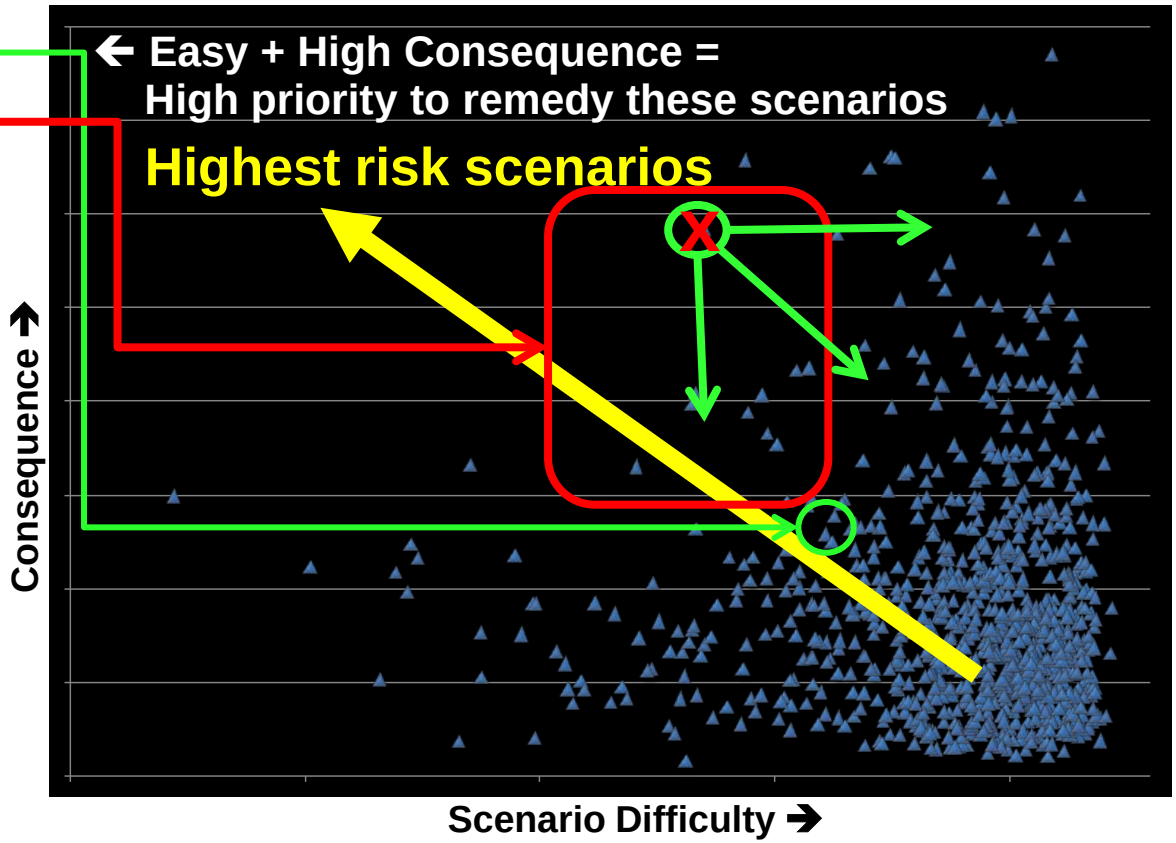

Parallels to Game Theory

Scenarios with the highest net utility are most advantageous, and most likely to be selected.

$$\frac{\Sigma\ Benefits \quad [\sim Consequence]}{-\ \Sigma\ Costs \quad [\sim Difficulty]}$$
$$Net\ Utility$$

This representation of security risk can be used for game theoretic assessments of attack scenario likelihood!

To "fix" a scenario we must

- **Eliminate it (make it impossible to achieve)**
- **Reduce the consequences that occur if it is completed**
- **Make it harder to accomplish successfully**

… or any combination of these

Sandia National Laboratories

A Notional Example Application

How do we decide which vulnerabilities should be addressed first?

- **Generally, work on scenarios that are both easy to do & high consequence.**
- **Enterprise decisions may be affected by intelligence data**
- **Decision maker values affect whether [Easy, ↓C] or [Hard, ↑C] is next**

Why use scenario difficulty as a component of risk?

- **Difficulty better reflects adversary planning processes**
- **Difficulty changes more slowly and predictably than likelihood**
- **Problem: How do we quantify the difficulty of an attack?**
  - ***This is the subject of ongoing research…***

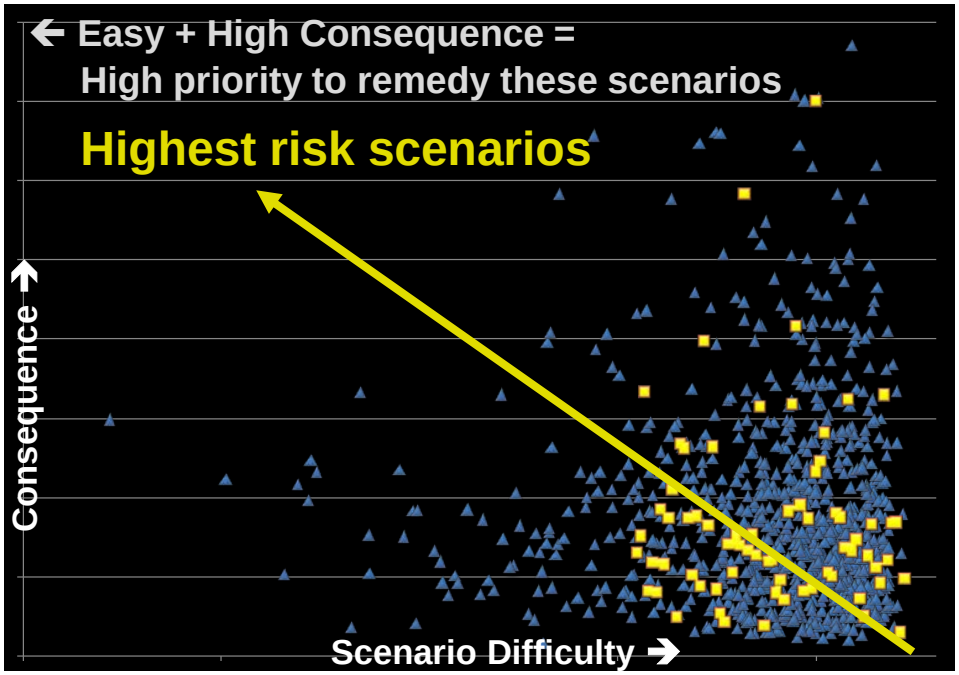

Composite (Enterprise/Facility) View of Security Risk

Investment insights from this method seem more robust & defensible than those based on highly uncertain attack likelihood estimates.

Considerations for Estimating Attack Scenario Difficulty

LDRD
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

Attack Preparation

- *Outsider attack participants*
  - *Number of engaged participants*
  - *Training & expertise required*
- *Insider attack participants*
  - *Number and coordination*
  - *Level of physical and cyber access required, sensitivity, vs. security controls*
- *Organizational support structure*
  - *Size, capabilities & commitment*
  - *Training facilities, R&D, safe haven, intelligence & OPSEC capabilities…*
- *Availability of required tools*
  - *Rarity, signatures for intelligence or law enforcement, training signatures…*

Attack Execution

- *Ingenuity & inventiveness*
- *Situational understanding*
  - *Observability & transience of vulnerabilities*
- *Stealth & covertness*
- *Dedication & commitment of participants*
  - *Risk to both outsiders & insiders includes personal risk, willingness to die, etc.*
  - *Risk to the "cause" or support base*
- *Operational complexity/flexibility*
  - *Precision coordination of disparate tasks*
  - *Multi-modal attack (cyber+physical+???)*

Scenario difficulty is a property of the *target.*
It estimates how capable the adversary must be to have a successful attack.

Risk managers can then ask, "Are the easiest attacks difficult enough to deter the adversaries we are concerned about?"

Estimating Difficulty of Attack Scenarios

General characteristics used to establish levels of difficulty for dimensions.

Level 1	Level 2	Level 3	Level 4	Level 5
Easy to get/do	Moderately easy to get/do	Difficult	Very difficult	Extremely difficult to get / do
Capability available by legal means	Requires capability similar to criminal activity	Requires capability similar to organized criminal activity	Requires sophisticated capability similar to large corporation	Requires state-supported capability
Requires no special skills	Requires low-level skills (~days of training)	Requires moderate-level skills (~months of training)	Requires high-level skills (~years of training)	Requires highly specialized skills (~multiple years of training, such as an advanced degree)
Easily accessible by general public	Accessible by public that has moderate-level knowledge	Typically accessible by criminal, paramilitary, or terrorist enterprises	Accessible by highly specialized organizations	Typically accessible only by elite forces
Essentially no early warning signatures - little risk to adversary of disruption	Some early warning signatures that may elevate general concerns of authorities – some risk of disruption			Very large early warning signatures – great risk of disruption

Example Scenario: Oklahoma City Bombing

Scenario 3: Oklahoma City Bombing. This scenario reflects the difficulty that was likely encountered by the participants in the plot to bomb the Murrah Federal Building in Oklahoma City.

Level ***(Score)*** *[1, 2, 3, 4, 5 → 1, 3, 9, 27, 81]*

Attack Planning & Preparation	Participants	2 *(3)*	Several (~2-5); Small team
	Training	2 *(3)*	Self-taught; Open source info; No professional foundation; Practice not required for critical tasks
	Support	1 *(1)*	Minimal; Few if any support personnel / collaborators; No intelligence support; Preparations easily concealed—no need for cover; Open source info
	Tools	2 *(3)*	Legal availability controlled, limited to special purpose uses; Typical of criminal enterprises
	# of Insiders	1 *(1)*	None
	Insider Access	1 *(1)*	None
	Ingenuity	1 *(1)*	Very predictable, straightforward approach; Easily conceivable by knowledgeable public; Defenses likely to be well prepared / trained against
Attack Execution	Situational Understanding	1 *(1)*	Minimal; Requires little recognition or utilization of exploitable conditions; Exploitable vulnerabilities are persistent and predictable, with evident signatures
	Stealth & Covertness	1 *(1)*	Minimal
	Outsider Commitment	2 *(3)*	Persistent remote exposure or participants, limited direct exposure to less-than-lethal conditions; Little risk of casualties, but significant risk of participant attribution
	Insider Commitment	1 *(1)*	None
	Complexity	1 *(1)*	Single avenue of attack with simple tasks; Unimodal tasks; If multi-modal attack, modalities are sequential, temporally decoupled
	Flexibility	1 *(1)*	Singular binary course of action; No contingency planning; Little tactical adjustment
Aggregated Score		*(21)*	*Score for each level is 3x that of the next lower level in this example.*

Summary

Risk-informed security investment prioritization is possible *if* risk is based on scenario difficulty.

- **Robust against likelihood uncertainties that constrain today's risk-based security decision-making.**
- **Difficulty reflects known adversary planning process better than likelihood.**
- **Communicates well with decision makers even if it cannot be used to roll up risk into a single number.**