Risk Informed Management of Enterprise Security
(RIMES)

or, "Why computing security risk based on a probability of attack is possible, but is likely not useful for Risk Management."

Presented at PSAM 16 – June 27-July 1, 2022 – Honolulu, HI

By Gregory D. Wyss
Distinguished Member of Technical Staff
Sandia National Laboratories

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

Sandia National Laboratories is a multi-program laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC (NTESS), for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
This document was approved for unlimited release as SAND2022-8311 C.

Hidden Dependencies in the Probability of Attack

- **To estimate P(Attack) one must estimate the likelihood…**
  - … that some known or unknown individual or group will exist, …
  - … during some specified time period, and will …
  - … decide that an attack can achieve an outcome [consequence] they desire, ***and*** …
  - … understand and validate an exploitable vulnerability or pathway to plan an attack, ***and*** …
  - … obtain the weapons, tools, skills and information required to accomplish the attack, ***and*** …
  - … decide that the attack's likelihood of failure, potential losses and risks are acceptable, ***and*** …
  - … decide that the costs and sacrifices required to accomplish the attack are acceptable, ***and*** …
  - … decide that this is the best opportunity to accomplish a desired objective by comparison to all other known opportunities at *this or any other* facility or location.

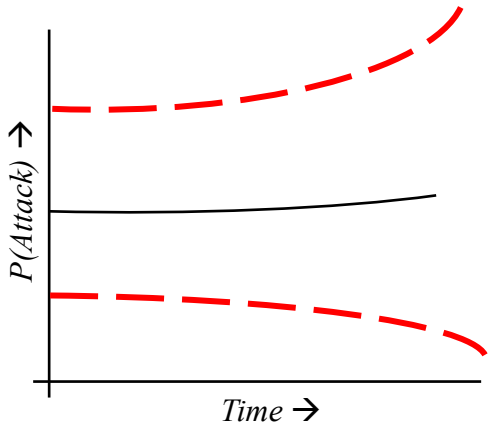

- **Uncertainty in P(Attack) should consider:**
  - What might cause new adversaries to exist, or old ones to cease? *World events? Personal events?*
  - How and why might adversaries' value sets change?
  - How and why might the adversaries' opportunities, required resources and tasks, or exploitable vulnerabilities change? *New technology? Changes in facility?*
  - How are all of these changes related to time?

Security Risk Management Assertions

- **Major, high-value security risk management decisions are being made all the time – often without meaningful risk information.**
- **Many problems exist in traditional security risk management tools** (e.g., a "Design Basis Threat").
  - Misleading "quantification" of security risk
  - Security discussions that focus on the wrong things
  - Bad risk management decisions
- **Quantification of security risk via PRA frequently makes these decisions less clear, not more so.**
- **Probability of attack leads to risk management problems for extreme-consequence attacks.**
  - Risk values and risk management are *extraordinarily uncertain* when probability of attack is elicited
  - A security risk method is essentially useless without addressing deterrence and threat shifting
  - This is especially true when "probability of attack" is elicited from expert judgment
- **Game theory is ill-equipped to address these issues for multiple adversaries vs. multiple targets.**
  - Practical and mathematical issues abound – multi-player game theory does not apply
  - Particularly ill-equipped for rare attacks that can lead to extreme consequences, due to uncertainties

Safety vs. Security Risk

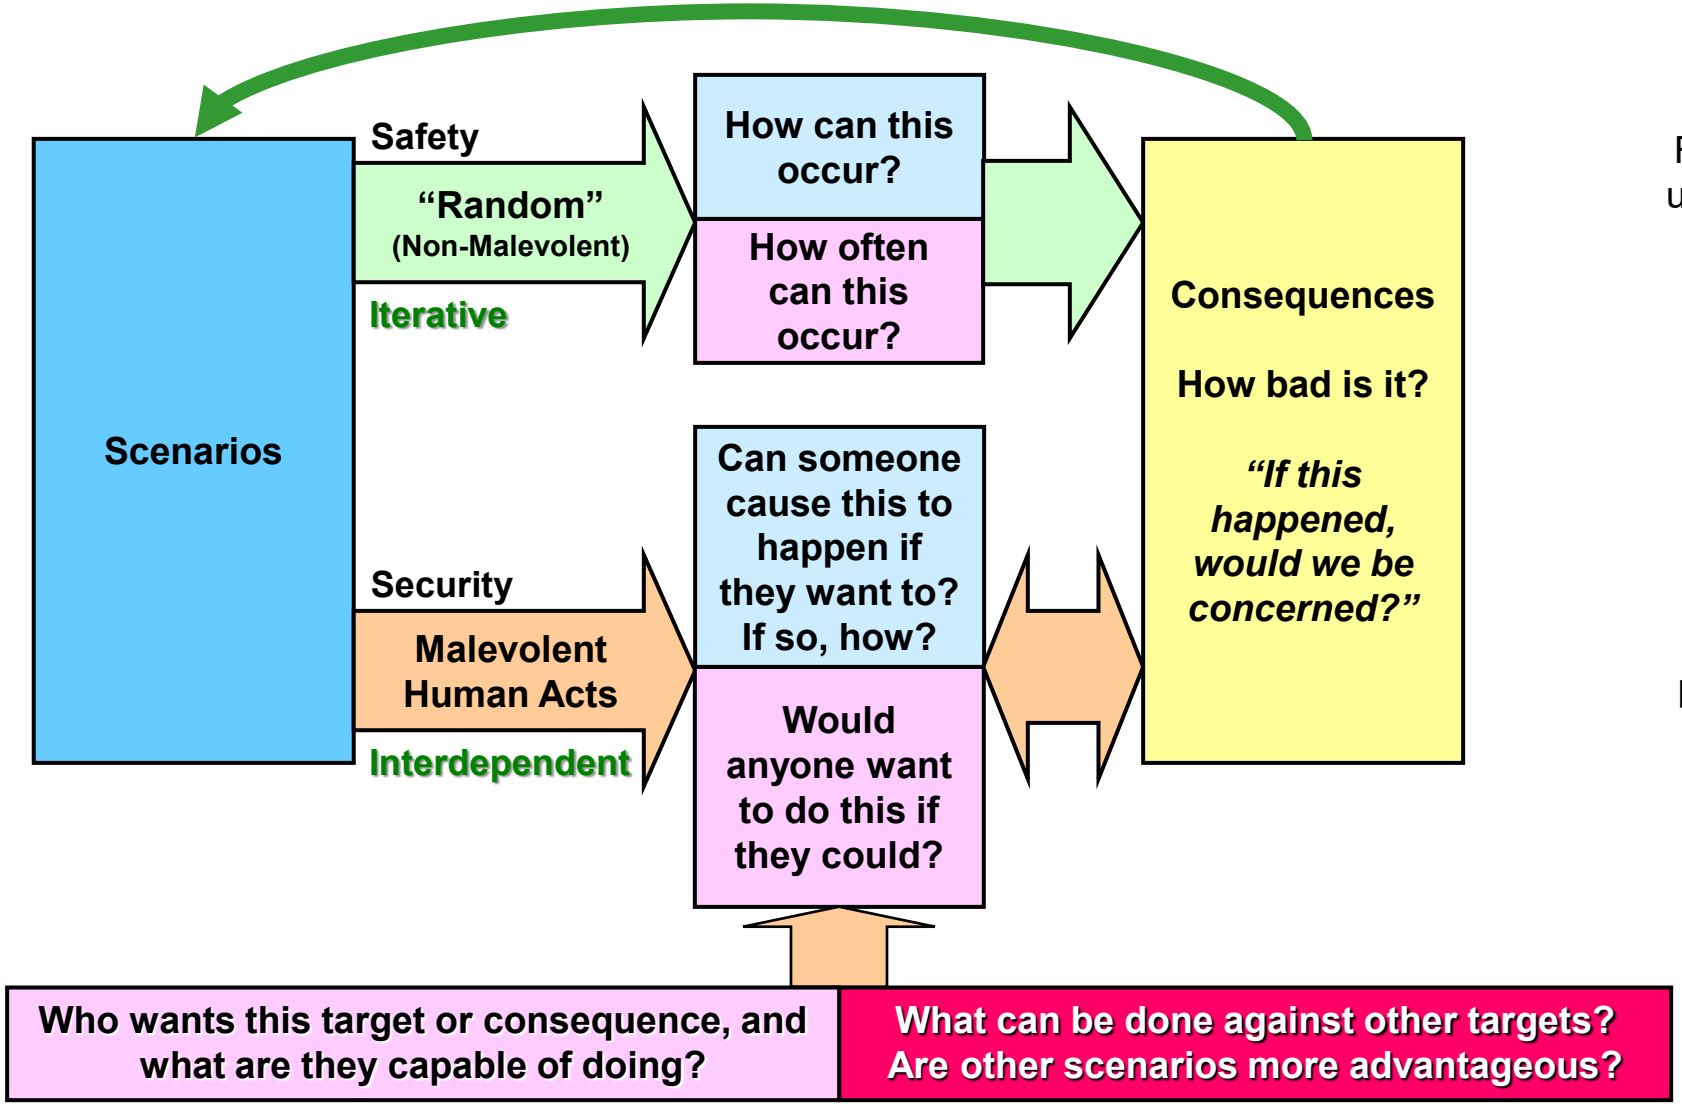

Results

Risk

Risk is the potential for realization of unwanted, negative consequences of an event

Risk Assessment

Systematic process to comprehend the nature of risk, express and evaluate risk, with the available knowledge.

Risk Management

Activities to handle risk such as prevention, mitigation, adaptation or sharing. It often includes trade-offs between costs and benefits of risk reduction and choice of a level of tolerable risk.

From the Approved Lexicon of the Society for Risk Analysis

Sandia National Laboratories

Difficulties with Probability of Attack

1. **Conditionality of Terms in the traditional safety risk equation risk equation does not hold up for security risk. All of the terms are interdependent.**

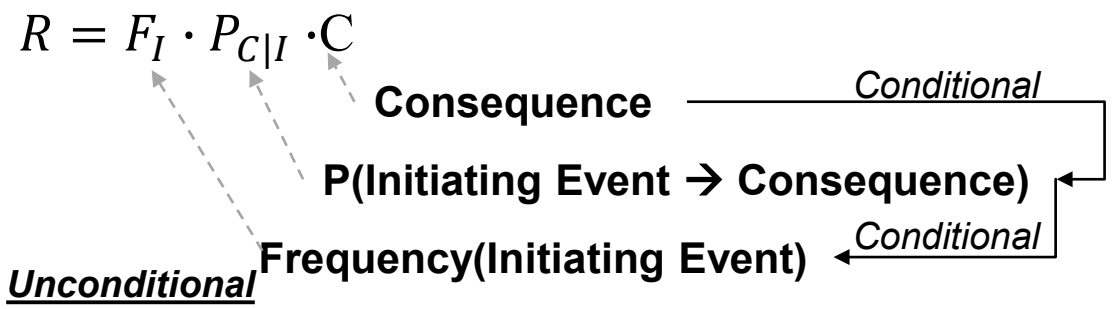

Safety Risk Equation

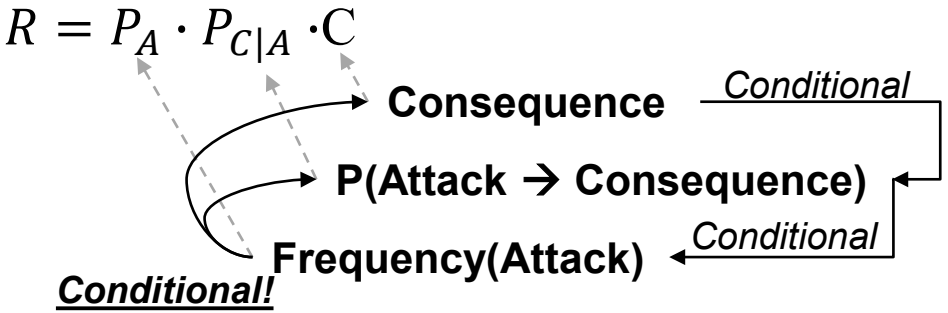

Security Risk Equation

2. **Norm Rasmussen said, "I do not believe that the safeguards [*i.e., security*] risks can be quantified using these [PRA] procedures" because P(attack) does not possess certain important statistical properties (e.g., randomness).**
   - 1976 remarks by Norm Rasmussen, MIT professor, "godfather" of PRA for nuclear industry, in: N.C. Rasmussen, "Probabilistic Risk Assessment: Its Possible Use in Safeguards Problems." Presented at the Institute for Nuclear Materials Management meeting, pp. 66-88, Fall 1976. *(see pg. 71)*

Sandia National Laboratories

Difficulties with Probability of Attack

3. **When conditionality is wrong, important characteristics of human behavior are not computable: Deterrence and Threat Shifting. These are critical to risk management.**
   - National Research Council. 2008. Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change. Washington, DC: The National Academies Press. https://doi.org/10.17226/12206.

4. **Long-term attack frequency has very broad uncertainties.**
   - **If not fully considered: Point estimates used for security risk values**
     - Decision makers can base their security risk management decisions on misinformation
   - **If fully considered: Uncertainty in security risk values is overwhelming**
     - Security risk management insights unactionable due to lack of statistical significance

These are fundamentally the same issues that were being debated when I entered the field of PRA over 30 years ago.

We need to focus on security risk management, rather than quantifying "how much or little risk exists."

Understanding and Managing Risk in Security Systems for the DOE Nuclear Weapons Complex (emphasis added) National Academy of Sciences, 2010, theme of Chapters 3 & 5.

Problems Managing Security Risk "By The Numbers"

Examples of change in system or situation:

- New adversary formed or discovered
- Adversary's values change
  - Desired outcomes
  - Risk tolerance: acceptable P(failure) or C(failure)
  - Budget, acceptability of costs
- Adversary's capabilities change
  - Weapons, tools, skills, information
- Availability of attack technology changes
  - New tech, or availability of existing tech
  - Cost to become an adversary who can successfully attack
- Information availability
  - Surveillance tools, Wikileaks, …
- Changes at this facility
  - Security upgrades, or new vulnerabilities introduced
- Changes at other facilities
  - May affect whether "my" facility is a desirable option for mounting an attack

Effect on Security Risk Management:

Adversaries' mindset is influenced ➔

P(attack) value may change ➔

Security Risk values may change ➔

Risk Landscape may change ➔

Mitigation metrics may change ➔

Need to redo the analysis?

Using P(Attack), security risk estimates can be fragile
(in addition to highly uncertain)

The mathematical model used in quantitative probabilistic security risk management cannot represent essential elements of security risk, except by re-eliciting P(Attack).

Adversary Adaptation

Deterrence

Threat Shifting

Quantitative security risk values can change dramatically overnight.

This makes them inappropriate metrics for long-term risk mitigation decisions

What were we trying to solve with RIMES?

- **Fundamentally, RIMES is a method for *long-term* security risk management without using P(attack).**
  - **Risk *management* without numerical risk *computation***
  - ***Long-term* (strategic) timeframe where P(attack) is highly uncertain and may change dramatically because of unknown/unpredictable events**

- ***Short-term* (tactical) security *relies* on intelligence-informed P(attack) to rapidly deploy protective measures to identified targets.**
  - **P(attack) – even qualitative – is meaningful and useful for this!**

RIMES Goal: Manage Security Risks

- **Problem: attack likelihoods are highly uncertain and change rapidly.**
  - **Depends on attacker's capability, motivation & intent**
  - **Depends on attacker's other opportunities inside *and* outside the system.**
  - **Predicting likelihood makes risk hard to use for security decision making**

- **A different risk management 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**

Attack scenarios:

Easy

&

High-Consequence

=

High Risk

These are the things that a defender can control and change.

Security Risk Management: Making Easiest Attacks More Difficult

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

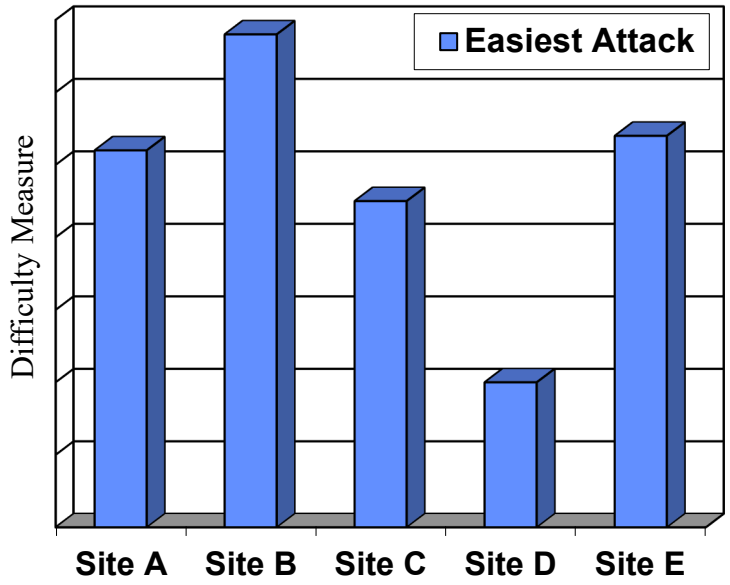

- Are sites balanced?
- Where should I spend my next dollar?

- How much have I improved?
- Why do my sites not meet the new security goal?

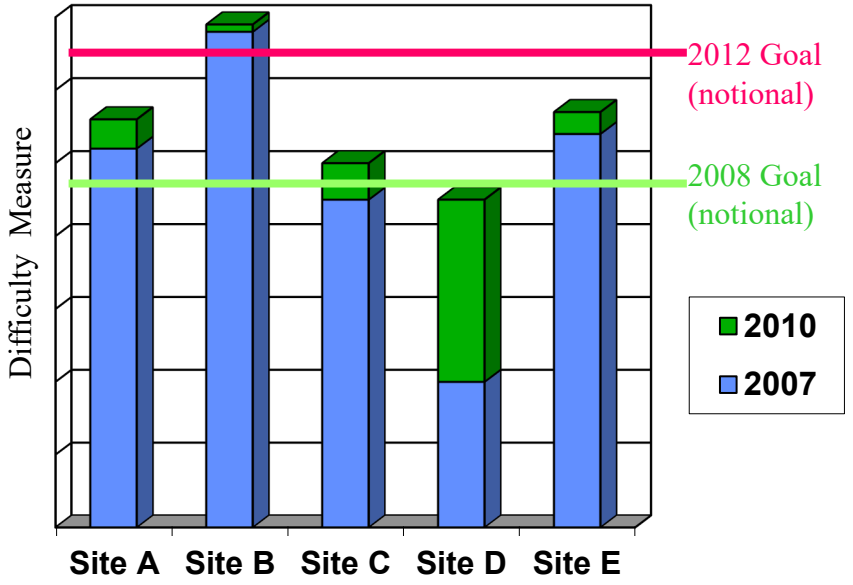

The Next Step: Manage Risk with Both Scenario Difficulty *and* Consequence

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!

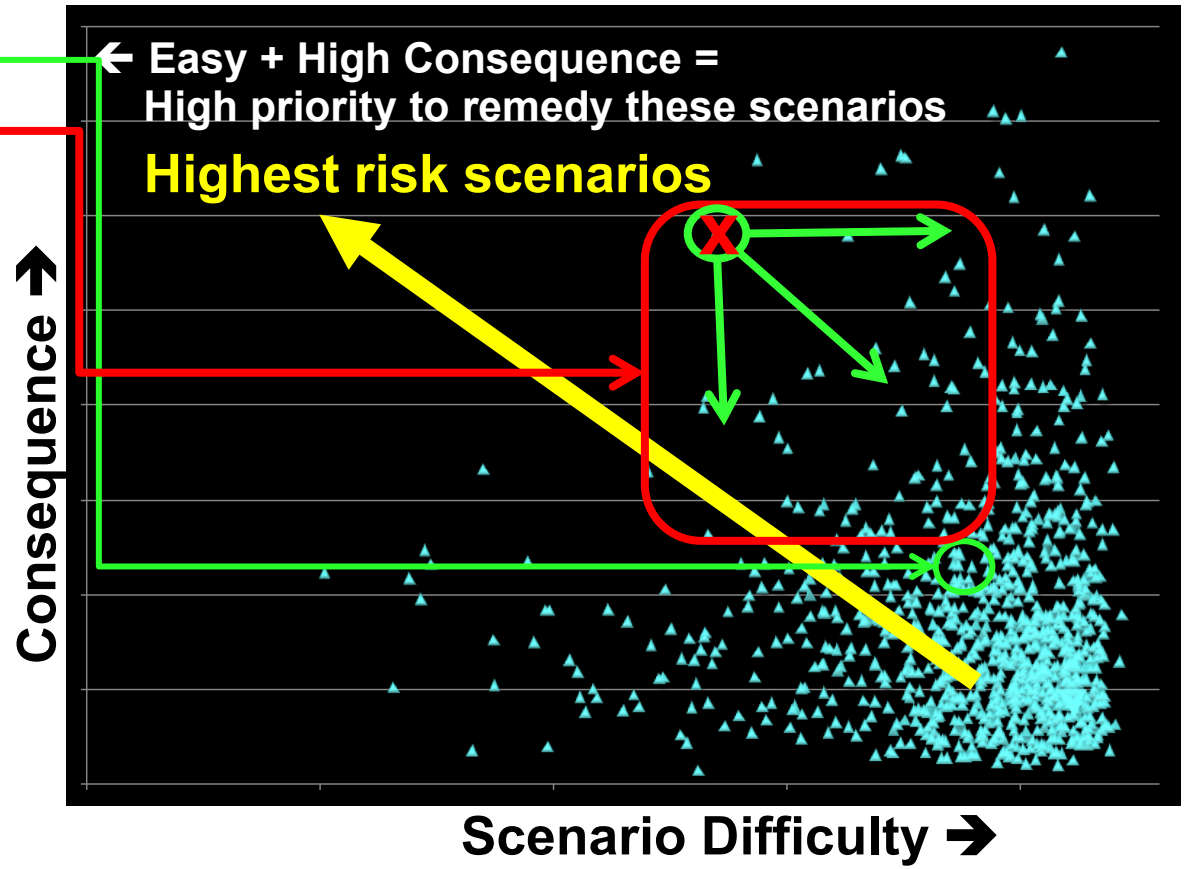

Why use scenario difficulty in security risk management?

- Difficulty better reflects the adversary planning process
- Difficulty changes more slowly and predictably than likelihood
- We have developed a qualitative (semi-quantitative) method to rank attack scenario difficulty

To "fix" a scenario we must

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

… or any combination of these

Considerations for Estimating Attack Scenario Difficulty

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?"

Quantifying Security Risk

- Can we quantify security risk? YES
  - It is always possible to quantify our understanding using Bayesian methods
  - For rare events driven by human choice, uncertainties are very large
- Is quantified security risk useful? Maybe
  - Broad comparison of disparate risks… may be useful
  - Broadly compare safety and security risks... may be useful
  - Detailed comparison of similar risks... maybe not...
    - Otherwise-clear risk mitigation decisions may be clouded by broad uncertainties – which are introduced by Bayesian estimation of adversary decision processes

"The likelihood of an attack should be an *output from* a security risk analysis, *not an input to it.*"

- Anthony Cox, Former Editor, *Risk Analysis* (emphasis added)

Questions?