QUANTITATIVE ANALYSIS OF CYBER RISK
HOW DO WE BEST MANAGE IT?

SPACE FOUNDATION SYMPOSIUM
APRIL 16, 2018

Elisabeth Paté-Cornell

and her Engineering Risk Research Group

Management Science & Engineering

Stanford University

Bases of our risk analysis work

- Quantify uncertainties using probability, including human & organizational factors
- System's dynamics and adversarial games.
- Statistics when they are relevant and sufficient, scenario analysis otherwise
- Objective: provide the best information we can to a decision maker to set priorities

Our cyber risk research: 5 vignettes

1. Statistical analysis of a specific data base of attacks for a fictional “Space Corp.” (Kuypers)

2. Network analysis and optimal connectivity. Application to a “smart” grid. Data from Sacramento Municipal Utility District (Smith) .

3. Dynamic analysis of the optimum replacement schedule of OS software. Motivated by the (mis) management of a water distribution system. (Keller)

Current research

4. Warnings of attacks

At three stages: before intrusion, after penetration, and at time of exfiltration. Objective: to mitigate the damage (Isaac Faber)

5. Fake news

Risk, and effectiveness of warning: detections, and corrections of fake news. Focus on elections & national security (Travis Trammel)

Quantification of cyber risk

Mathematical approaches in 5 PhD's

Three ways to capture uncertainties in risk curves (probability of exceeding loss L)

1. A statistical analysis of data (if *relevant* ones exist)

1. A probabilistic analysis (scenario-based)

1. Both combined (on the tail of the loss distribution)

Elements of our cyber risk model for a specific organization

Target-specific information:

- The nature of the target organization
- The information to be protected
- The structure of the system (physical and cyber)
- The potential, most likely, adversaries
- The consequences of a successful attack

- Statistical data analysis when they exist
- Bayesian network to model potential attack scenarios that we have not seen yet.

Two distinct kinds of cyber attacks Example of "Space Corp."

- Operational, routine attacks on organizational systems, for which statistical may have been gathered (often, most of the cost of cyber risk)

- Catastrophic, destructive attacks that may not have happened yet but threaten the organization: requires in-depth analysis of attack scenarios

The distinction may be fuzzy (close calls)
but the data and the analyses are different

Focus first on daily operations: routine attacks and costs

➢ Types of attacks or accidents

- Lost or stolen devices
- Data spillage
- Email
- Website
- Malware

➢ Costs of a successful attack

- Investigation
- Direct costs
- Loss of privacy information
- Reputation damage
- Loss of intellectual and physical property
- Business interruption

Countermeasures

- Firewalls
- Full disk encryption
- Two-factor authentication (e.g., password, pin, etc.)
- System compartmentalization
- Data Loss (exfiltration) Protection
- Malware detection
- Email filtering
- Biometrics, etc.

Open DLP

Effectiveness of these measures

Depends (among other things) on

- The nature of the system attacked
- The type of attack (e.g., by insiders)
- The ease of implementation (16 character passwords?)
- The sophistication of the attackers

1. Empirical analysis of incident data
with Marshall Kuypers (based on statistics)

Data often exist but are well guarded. Here: 60,000 incidents over six years of various routine attacks (e.g., lost or stolen laptops) in a large organization

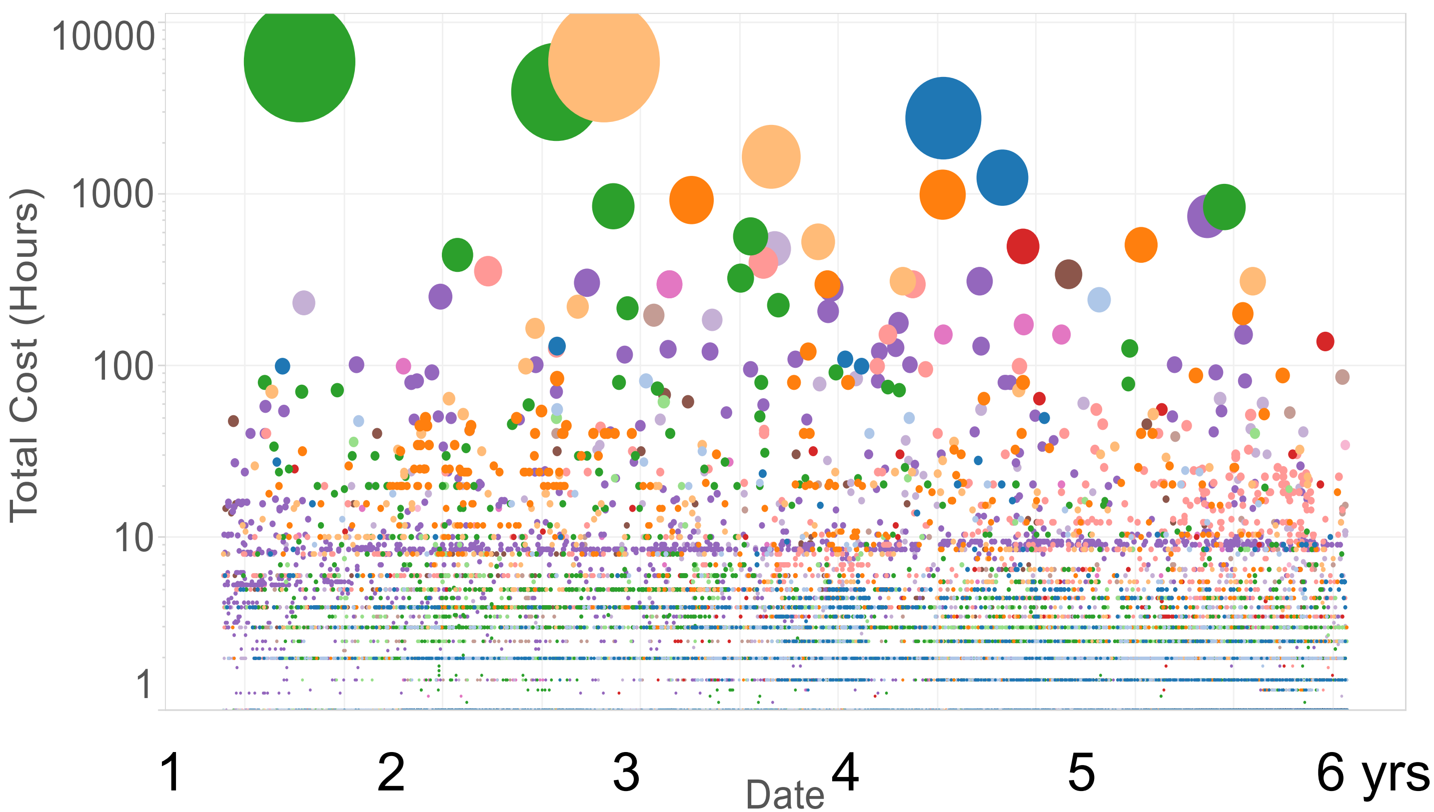

Statistical data and expert opinion to initialize probabilistic models ("Space Corp.")

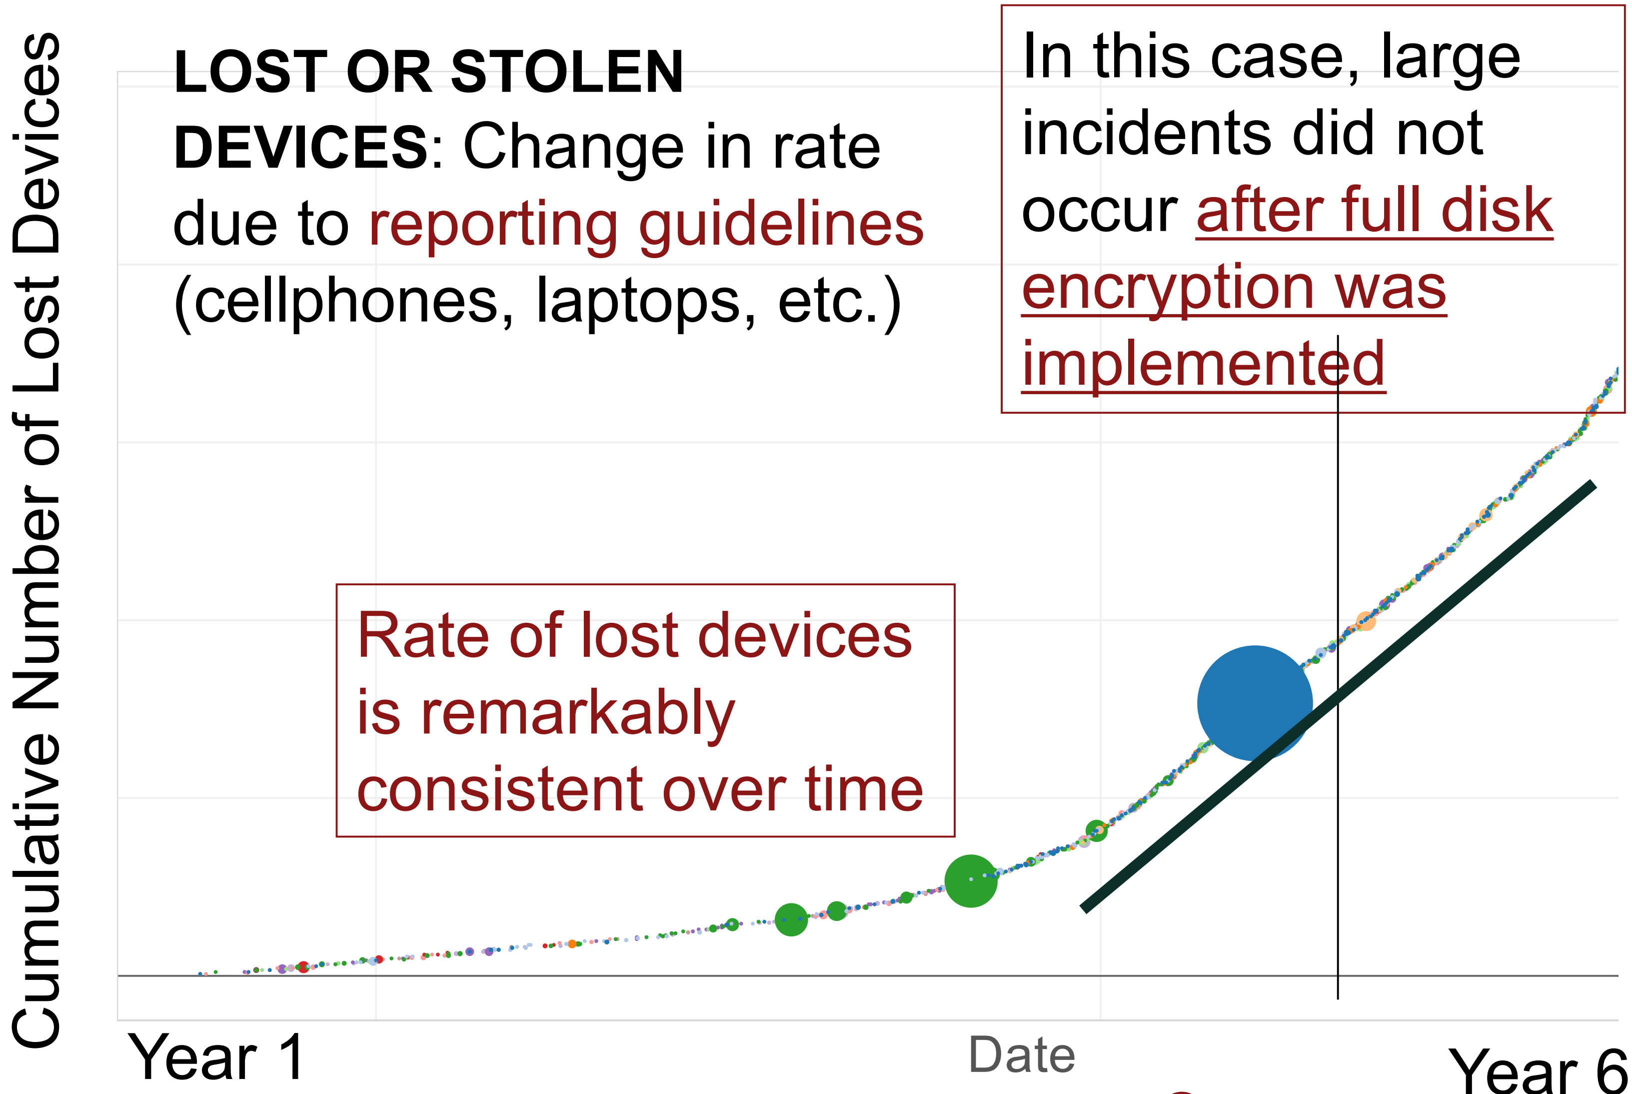

Combining statistical models with scenario analysis and probability

- Severe-impact incidents may already be included in the data.
- Large incidents that have not occurred yet require a scenario-based model (probabilities & losses)
- The two models overlap (e.g., close calls)
- Same cost analysis for both models.

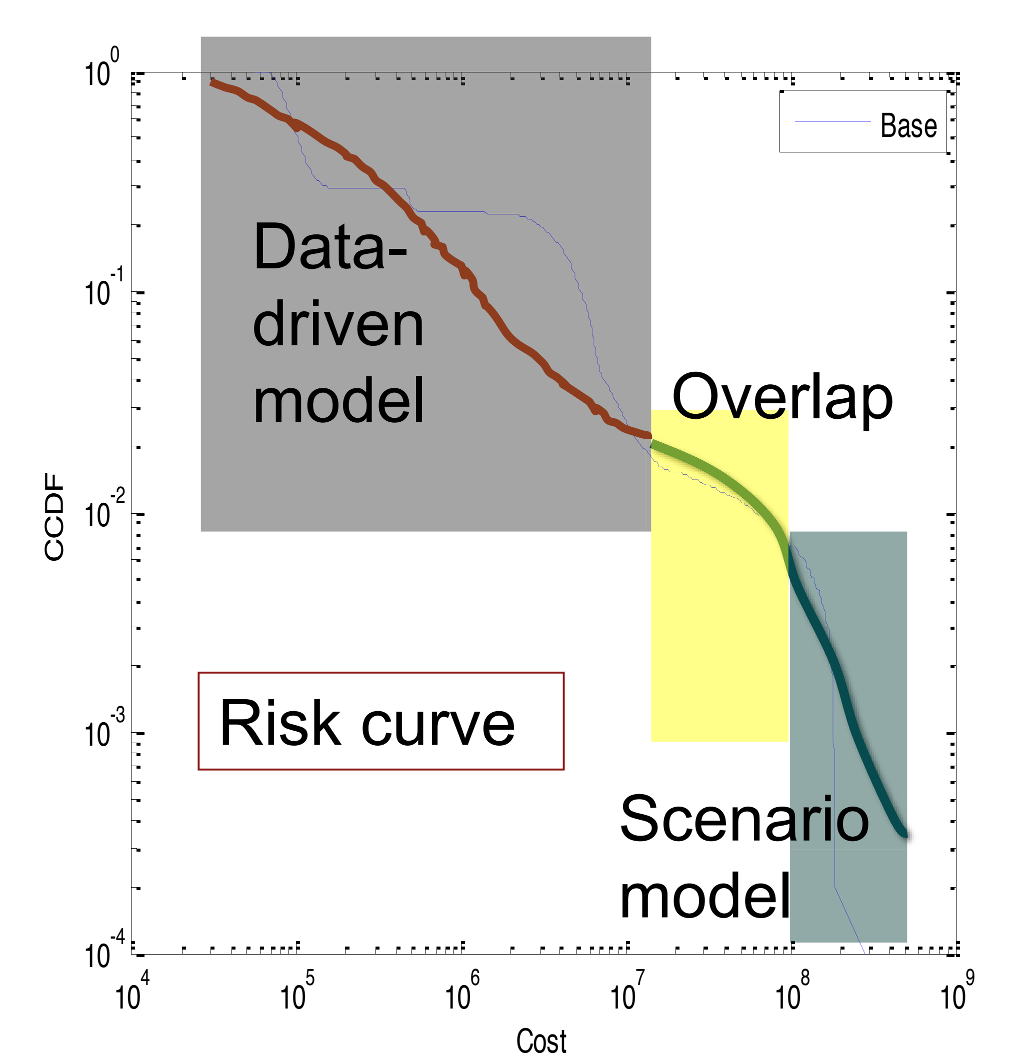

Takeaways

- Risk quantification can be done

  combination of *statistical analysis* (past attacks), and *future scenario analysis* (with probability) based on expert opinions and close calls

- Rate of attacks

  In this organization, relatively constant.

- Counter measures' effectiveness can be assessed and compared.

  In this case, Full-Disk Encryption and Two-Factor Authentication were showed to be most effective.

2. Network defense and optimal level of connectivity (with Matthew Smith)

- Smart Grid Benefits

  Adding communication *improves efficiency and reliability* by allowing grid systems and operators to react quickly to changing conditions (e.g., demand)

- But added connectivity increases vulnerability

  The smart grid is exposed to new digital threats: denial of service attacks, intellectual property theft, invasion of privacy, sabotage, etc.

The networks (physical and information) and possible cyber attacks

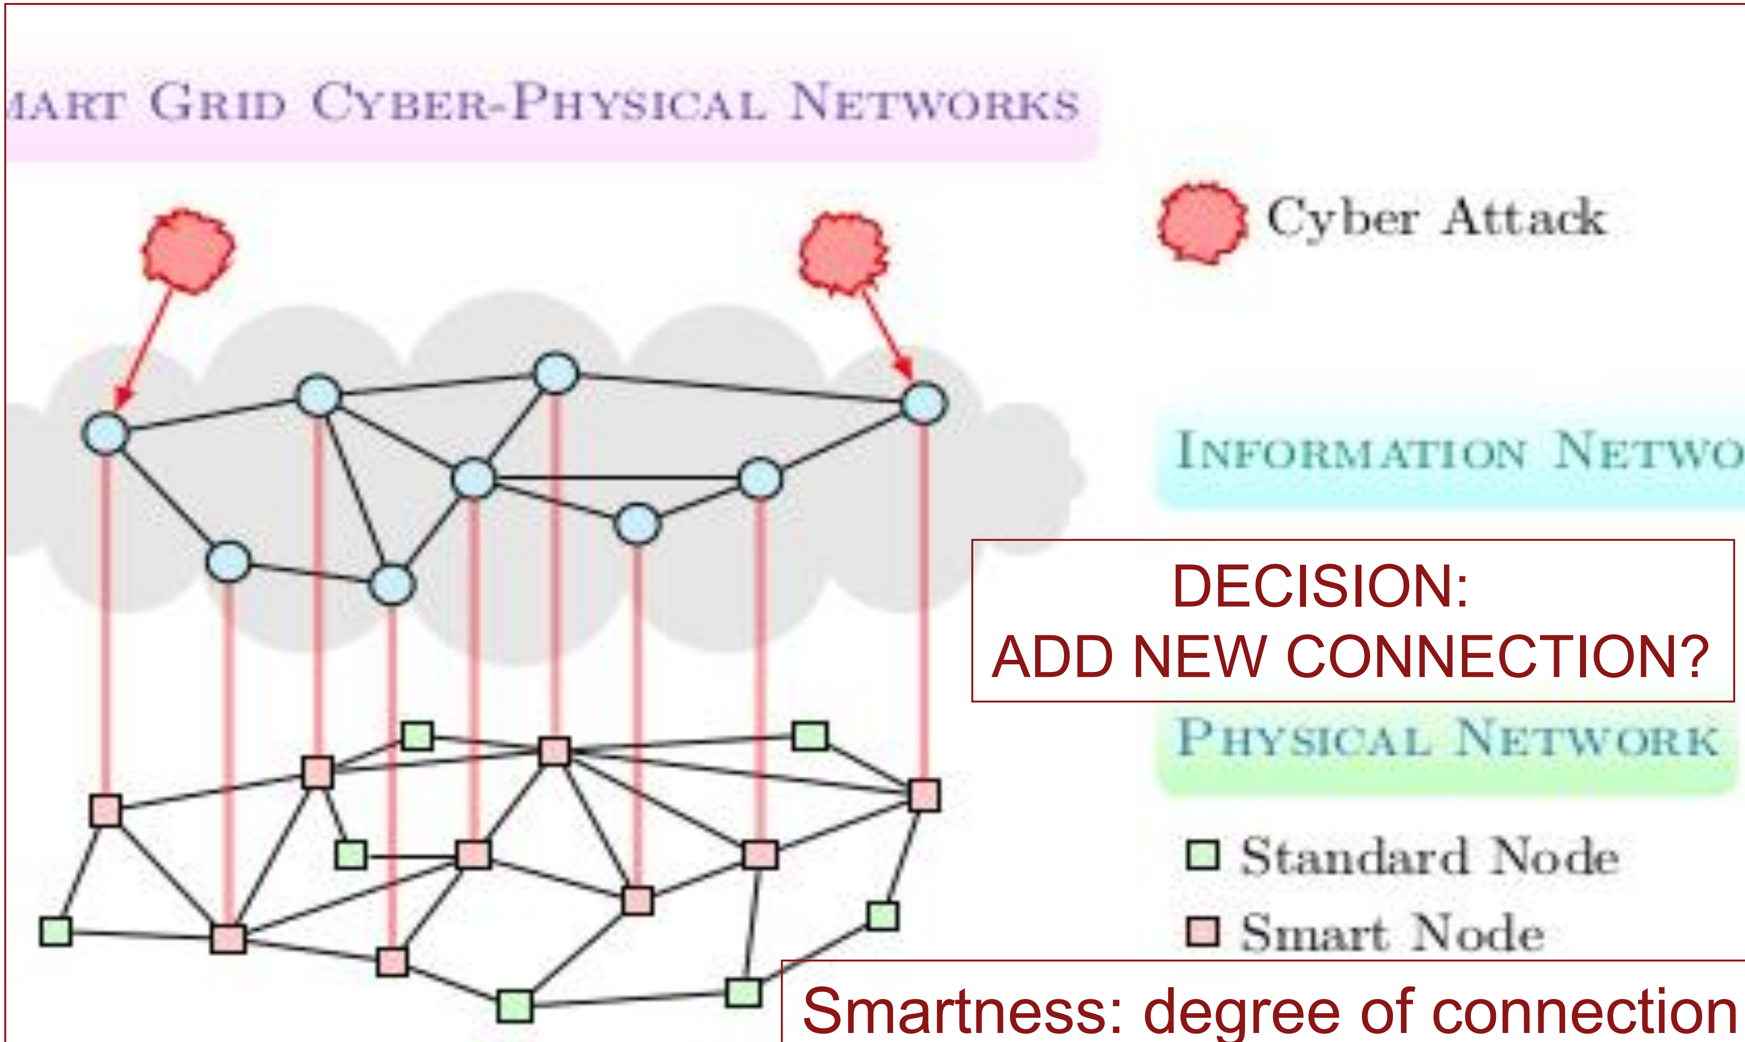

DECISION:
ADD NEW CONNECTION?

Smartness: degree of connection

Dynamics of Cyber Security Investments

- Focus here on proactive use of cyber defense teams for defensive and information gathering purposes
- Choice: Exploitation (of known vulnerabilities) vs exploration (find new ones). Classic Multi-Arm Bandit problem –o-> Multi-Node model

Defense

Information Gathering

US Department of Defense Cyber Protection Team

Search for Optimal Connectivity

Step 1 – Systems Analysis

Identify classes of cyber failure scenarios for a Smart Grid network based on structure

Step 2 – Economic Analysis

Evaluate financial benefit and risk of increased connectivity

Step 3 – Stochastic Modeling

Use Multi-Node Bandit security model to assess optimal protection against old and new vulnerabilities

Step 4 – Decision Analysis

Find optimal smartness, to support decisions of system operators

Results: optimal point where marginal benefit equals marginal risk

"Smartness"= the degree to which the physical network has been integrated into the information network (0 to1)

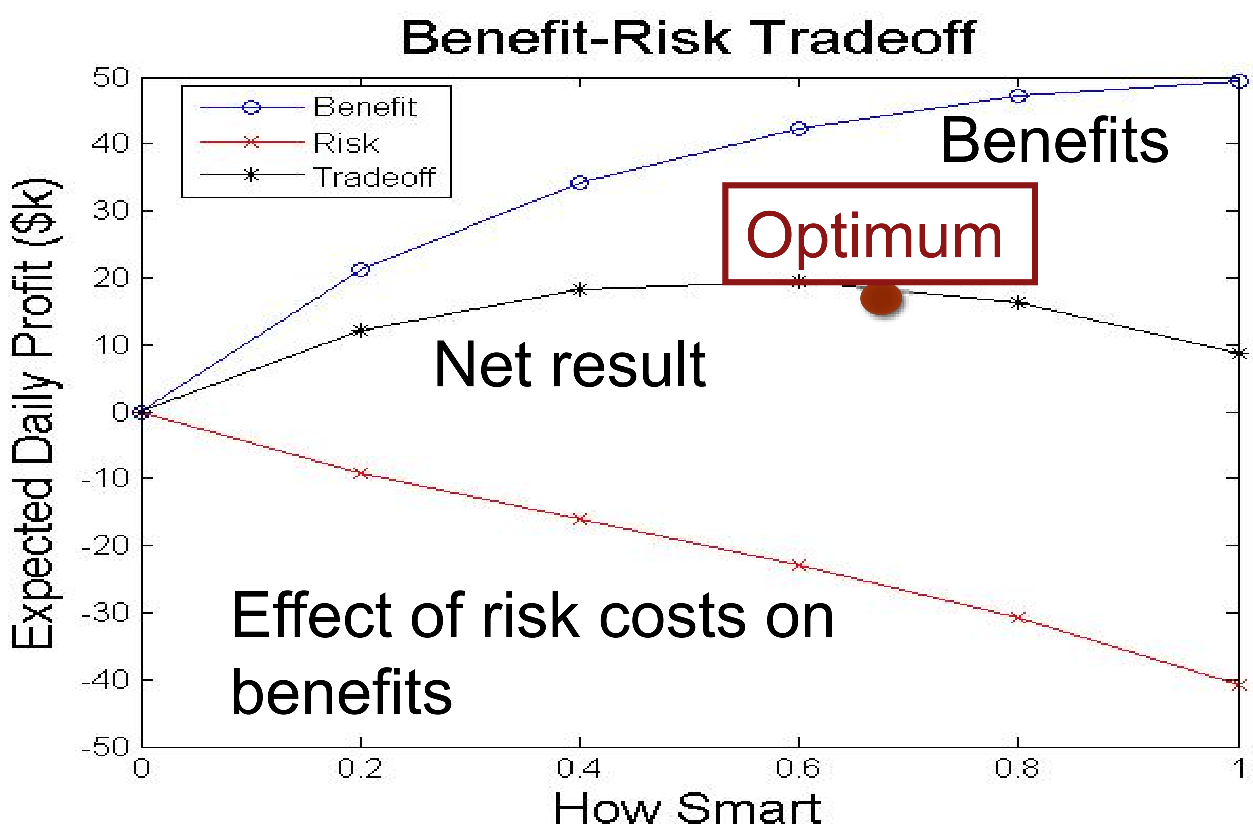

Takeaways

- "Smartness" in the electrical grid is beneficial up to a point.
- Risk management includes allocation of defense teams (exploitation vs. exploration).
- Optimum connectivity can be assessed through risk analysis (statistics and experts opinion).
- The first task is to understand the structure of the network and the potential for cascading effects given the interconnections.

3. Upgrading control software to stay ahead of an adversary with Philip Keller

- How often to upgrade the system?
  - New software or reconfiguring existing software regularly can complicate cyber attacks, at a cost
  - Ex. of a water distribution system (no upgrade for 10 years!). Same problem for hospitals.

- Examples of failures to upgrade operating software
  - The ransomeware attack of May 12 2017
  - The Ukraine electric hack: 6 months of surveillance

Dynamic system analysis

- Questions:
  - How long will it take to an adversary to penetrate the system and find the critical target? (random variable)
  - How often should the software be changed given experience, potential attackers, new signals and new malware?

- Factors involved in that decision:
  - Discovery of new software vulnerabilities
  - Software installation and infrastructure costs

- Illustration: water distribution system (attack after 10 years of no updating)

Reconfiguration and Patch Decisions

- Probability of successful attack for different system ages derived from existing data (from Symantec). As one waits:
  - Vulnerabilities accumulate
  - The adversary has more time for reconnaissance
- Decision analysis: combining probability and costs of a successful attack with costs of software changes

attacker/defender model=> optimum upgrading

1. Game Analysis: Model of adversary
2. Decision of the Malware Developer
3. Stochastic Model of Software & Patch Development
4. Stochastic Models of Vulnerability Discovery
5. Stochastic Model of Conflict
6. Result: optimum upgrading time

Costs and Results

- Costs:
  - Successful attack to the infrastructure; for example, lost productivity, or people without water
  - Down-time during software installation, and subsequent adaptation
  - Software licenses
- Result:

  Optimal timing of software replacement, or patch installation after release

Takeaways

- Need to change the software to stay ahead of an attacker trying to find its way into the system
- Optimum time determined by the speed of the attackers' progress, the emergence of new vulnerabilities or the resolution of existing ones
- Stochastic models (here, Markov) allow representation of the variation of the risk as time passes, and support of the decision to upgrade or change the defenders' software

4. Early Warning Systems for Cyber Security

with Isaac Faber

CURRENT RESEARCH OBJECTIVES AND METHODS

- Machine learning techniques for early stage attack to move ahead of damaging events
- Global honeypot sensor array to collect real data
- Communication system on changing risk profiles to issue warning for a given cyber system
- Use of industry standard attack graph, e.g., kill chains (attackers' plans): reconnaissance, weaponization, delivery, exploitation, installation, command and control, actions on objectives

Timeline: Example of Malware Attack

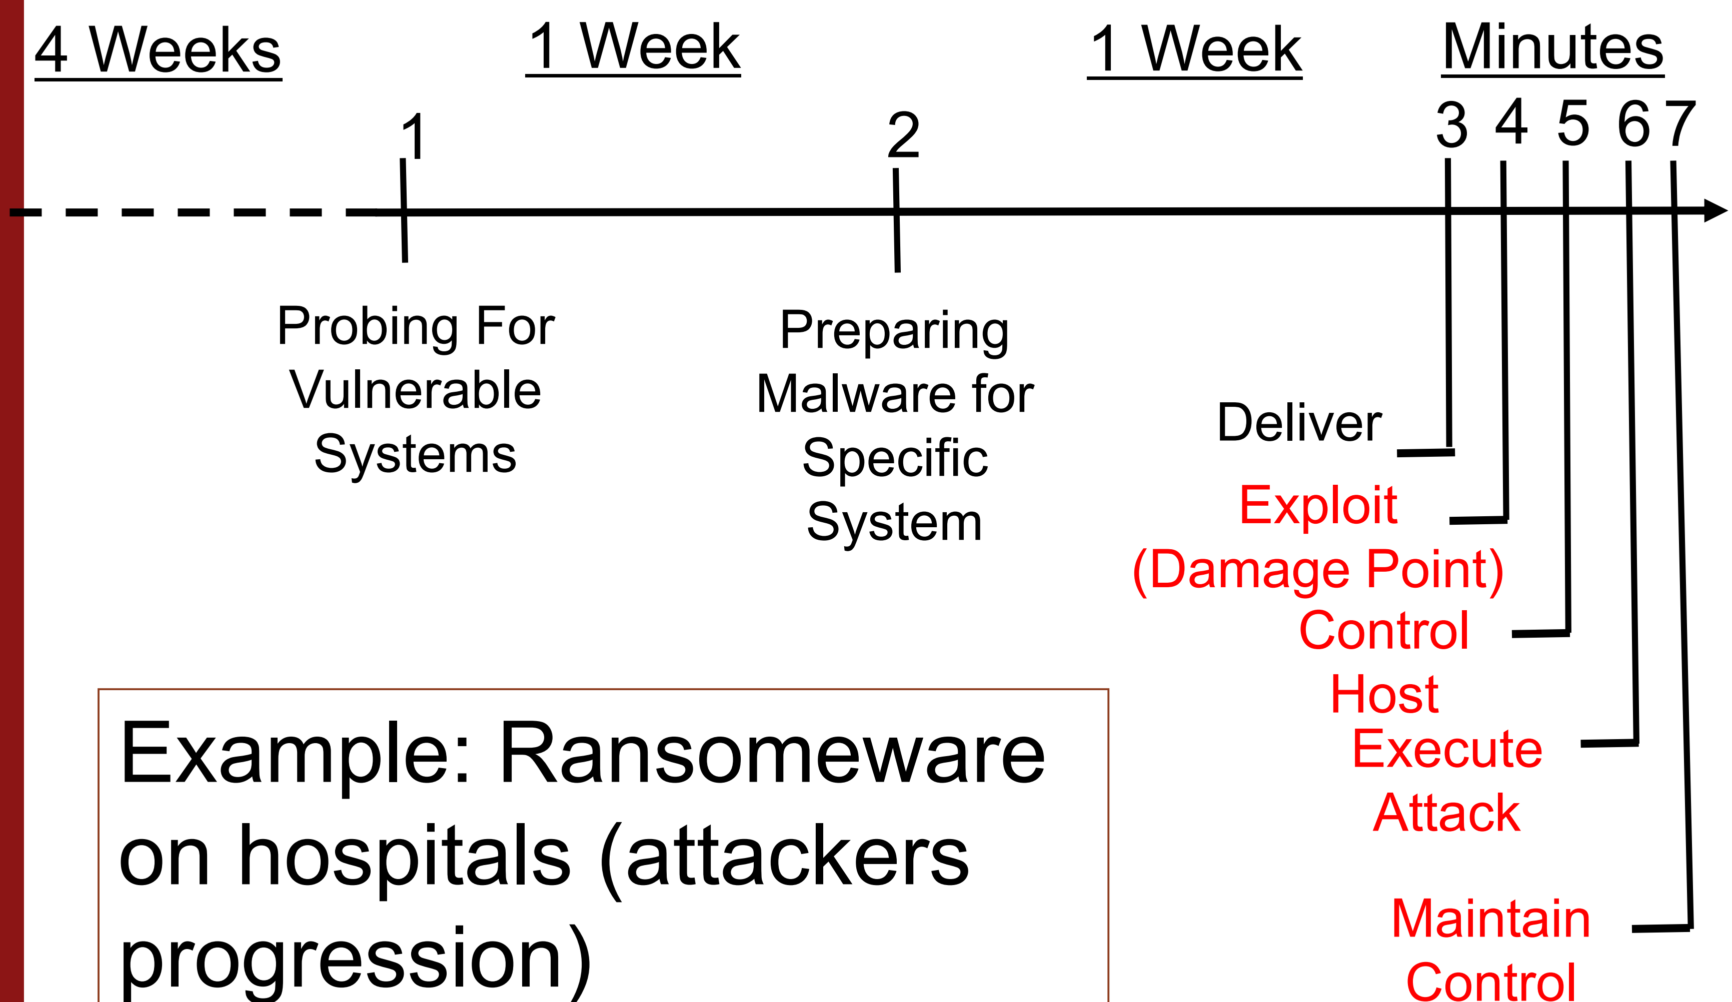

Example: Ransomeware on hospitals (attackers progression)

Honey pots: Locations and cloud service providers

Locations:

Virginia , USA

London UK

Toronto, Canada

Brazil

Frankfurt, DE

Seoul, South Korea

California, USA

Frankfurt, DE

South East Asia

Service providers

Azure

Amazon Web Services

Digital Ocean

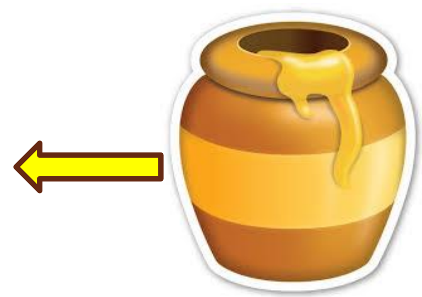

Computations

- Probability distribution of time to attack given raw sensor signals
- Probability distribution of severity (costs) of attack
- Identification of defensive/offensive countermeasures and decision cycles
- Probabilities of time to attack and effectiveness of countermeasures

Preliminary take aways

- Precursors of cyber attacks and behaviors can be observed early in the game, providing warnings of cyber threats *with some probability*

- Machine learning techniques involving deep learning seem to provide promising tools.

5. Fake news

with Travis Trammel

ONGOING RESEARCH

- Problem: U.S. government budget and funding allocation to combat sponsored fake news?
- Focus
  - › Financial
  - › Political (elections) and military attacks
- Objectives

  Anticipate, recognize (various degrees of "fakeness"), and counteract fake news at the earliest possible stage, in a credible fashion
- Timing is critical

Political and military examples

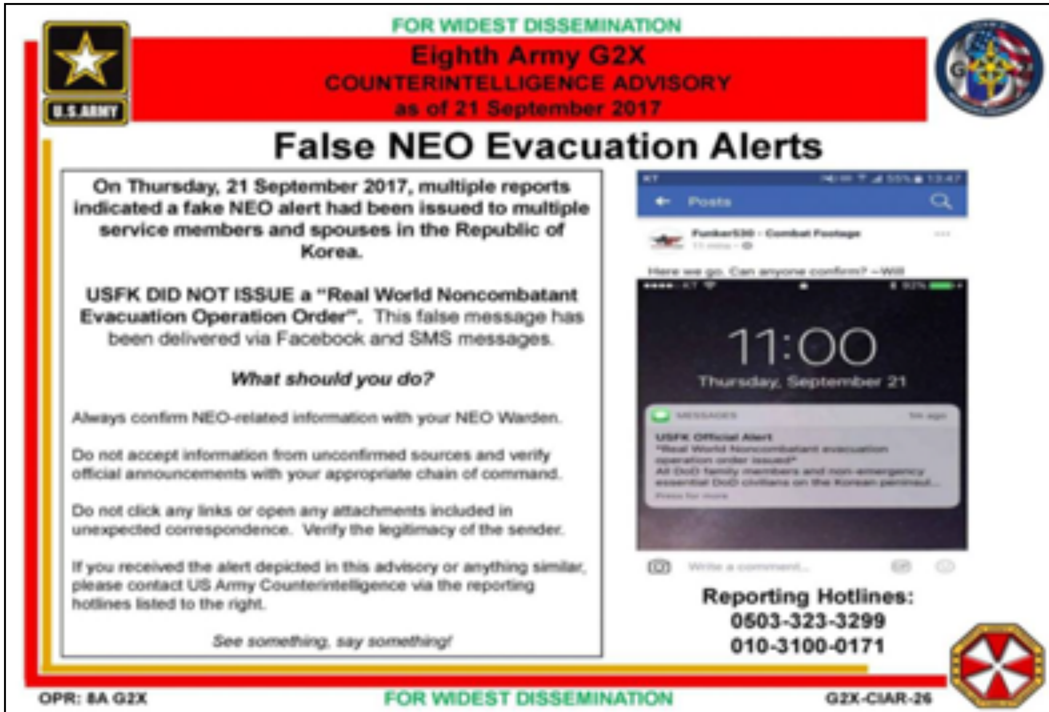

Fake evacuation alert of US military in Korea (2017)

Correction message

Russian false claim on NATO (04/2017)

Fake News Evolving Environment

- Connectivity and social media

  Vast amounts of information at unprecedented pace with global reach. Future global internet connectivity (51% global population connect in 2017)

- Technology (fake video & audio)

  Will make fake news more and more convincing. Ex: Russian use of a video game to simulate an American attack)

Probabilistic Risk Analysis and Adversary's Timing of Fake News

- Optimal timing of fake news by attacker if there is a targeted event (e.g., elections)
- => Anticipation by the defender

Example of elections

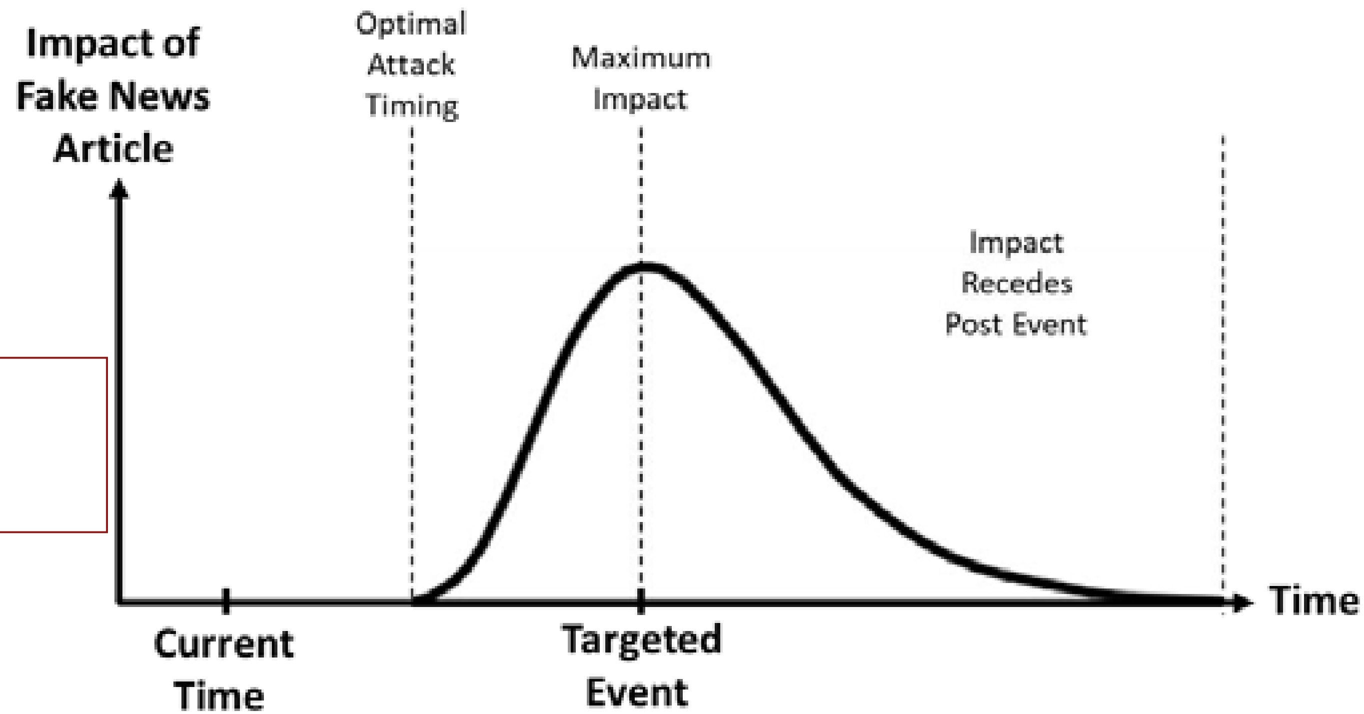

Countermeasures before and after attack

Some possible countermeasures:

Education

Flagging

Preliminary takeaways

1. There is a spectrum of fake news (and how fake) and probabilistic analysis allows assessing the chances of an attack's success
2. Detecting and correcting the obvious ones is step1.
3. Some can be anticipated (ex: elections in France)
4. The timing and the credibility of the response are essential to its effectiveness
5. Allocating resources may depend on the timing of the event of interest (e.g., elections) and on the geographic distribution of potential targets

Conclusions

The perception of cyber risk is often apocalyptic, but the real question is: what do we do next?

- There is a lot of qualitative research about the feasibility or legality of various protective measures.
- Accessing existing data sets and gathering new ones is key to the relevance of the results.
- Quantitative risk analysis is needed(and feasible) to bring some reality into perception and support rational decisions

A few years ago, cyber risk analysis was often deemed "impossible".

Now the question is:

How can we do it better?