Jens Grossklags[1], Benjamin Johnson[2], and Nicolas Christin[2]
[1]UC Berkeley, School of Information
[2]Carnegie Mellon, CyLab

The Price of Uncertainty in Security Games

Presented by Nicolas Christin at the Eight Workshop on the Economics of Information Security (WEIS 2009). University College London, June 2009.

Motivation

- Lack of good metrics to characterize judicious security investments
  - Marketing pitches vs. defensible metrics
  - Assessing penalties for cybercrime
- Economic models help, but usually assume full rationality and perfect information
- In practice:
  - Limited information due to size and complexity of network
  - Failure to discover optimal strategies
  - Failure to implement the chosen strategies

→ *How valuable is information in the context of security decision making?*

→ ***How do we even measure that?***

Contribution and Approach

- Propose and evaluate set of metrics to quantify value of information in information security decision-making

- Based on stylized network security games analysis
  - Under different information conditions
  - Under different expertise conditions

Background: Security Models

- Originally proposed in [GCC:WWW'08, GCC:EC'08] and presented at last year's WEIS

- Two key components of a security strategy
  - Self-protection (e.g., patching system vulnerabilities)
    - Joint protection level determined by all participants of a network
    - Public good
  - Self-insurance (e.g., having good backups)
    - Individual level of loss reduction
    - Private good

General Utility Model

Expected loss
(w/o countermeasures)

Network protection
level
(public good)

Insurance
purchased
$0 \le s_i \le 1$

Expected utility

$$U_i = M_i - p_i L_i (1 - s_i)(1 - H(e_i, e_{-i})) - b_i e_i - c_i s_i$$

Initial
endowment

Private
insurance level
(private good)

Protection
investment
$0 \le e_i \le 1$

Different contribution functions

- Weakest-link: $H(e_i, e_{-i}) = \min(e_i, e_{-i})$
  - Example: corporate network protection

  → $U_i = M_i - p_i L_i (1 - s_i)(1 - \min(e_i, e_{-i})) - b_i e_i - c_i s_i$

- Best shot: $H(e_i, e_{-i}) = \max(e_i, e_{-i})$
  - Example: Censorship resilient networks (see: Tor)

  → $U_i = M_i - p_i L_i (1 - s_i)(1 - \max(e_i, e_{-i})) - b_i e_i - c_i s_i$

- Total effort: $H(e_i, e_{-i}) = \frac{1}{N} \sum_i e_i$
  - Example: Peer-to-peer (swarming) transfers (see: BitTorrent)

  → $U_i = M_i - p_i L_i (1 - s_i)(1 - \frac{1}{N} \sum_k e_k) - b_i e_i - c_i s_i$

Uncertainty

- Expected losses may differ among players.
- Expected losses for other players may be unknown.
  - We assume that all expected losses are UID (uniformly and independently distributed) in [o,L].
- Some players may not take into account the expected losses of others.

Information Conditions

- Complete Information
  - You know all players' expected losses, including your own. E.g., (weakest link):
  - $U_i = M - p_i L(1 - \min_{j=1}^{N} e_j)(1 - s_i) - be_i - cs_i$
- Incomplete Information
  - You know you own expected loss but not others'. You know the distribution. E.g.,
  - $U_i = M - p_i L(1 - E(\min_{j=1}^{N} e_j))(1 - s_i) - be_i - cs_i$

A Mixed Economy

- *One expert player* acts strategically based on all available information.
- *All other players* choose levels of protection and insurance based on a straightforward cost-benefit analysis, ignoring behavior of others.
  - perceived utility:

$$U_i = M - p_j L(1 - e_j)(1 - s_j) - be_j - cs_j$$

  - actual utility:

$$U_i = M - p_j L(1 - \min_{k=1}^{N} e_k)(1 - s_j) - be_j - cs_j$$

Methodology

- For each information condition: complete and incomplete
  - Compute an expected utility for the expert player
  - Expert player's strategy: best-response to the behavior of the naive players.
- We take an additional expected value over all attack probabilities
  - Leave the final "expected utility" as a function of parameters known under incomplete information.

Price of Uncertainty

- Goal: measure how much uncertainty costs an expert player
  - Quantify a payoff differential between full information condition and limited information condition
  - Payoff depend on 5 parameters: initial endowment M, cost of protection *b* and cost of insurance *c*, number of players N, and magnitude of losses L
    - Need to reduce the number of parameters through the definition of the metric
- Three possible metrics
  - Difference metric
  - Payoff-ratio metric
  - Cost-ratio metric

Payoff Difference Metric

$$\max_{b,c\in[0,L]} [\text{Expected Payoff Complete}(b,c,L,L,N) - \text{Expected Payoff Incomplete}(b,c,L,L,N)]$$

- Worst-case difference in payoff between complete and incomplete information
  - Maximum taken over all possible prices for protection and insurance
- An insignificant price of uncertainty yields an output of zero
- The metric's output increases w/ the significance of the price of uncertainty

Payoff Ratio Metric

$$\max_{b,c\in[0,L]}\left[\frac{\text{Expected Payoff Complete}(b,c,L,L,N)}{\text{Expected Payoff Incomplete}(b,c,L,L,N)}\right]$$

- Somewhat analogous to "price of anarchy"
  - payoff-ratio of a game's socially optimal equilibrium to its worst case Nash equilibrium
- Currency independent
- An insignificant price of uncertainty yields an output of one
- The metric's output increases w/ the significance of the price of uncertainty

Cost Ratio Metric

$$\min_{b,c\in[0,L]}\left[\frac{\text{Expected Payoff Complete}(b,c,L,0,N)}{\text{Expected Payoff Incomplete}(b,c,L,0,N)}\right]$$

- Similar to the payoff-ratio metric, but with a different canonical choice of zero for the initial endowment M
  - Simpler algebraic analysis due to an abundance of term cancellations
- An insignificant price of uncertainty yields an output of one
- The metric's output *decreases to zero* w/ the significance of the price of uncertainty

Best Shot, Payoff Difference

Best-shot: Payoff difference as a function of number of players N and losses L

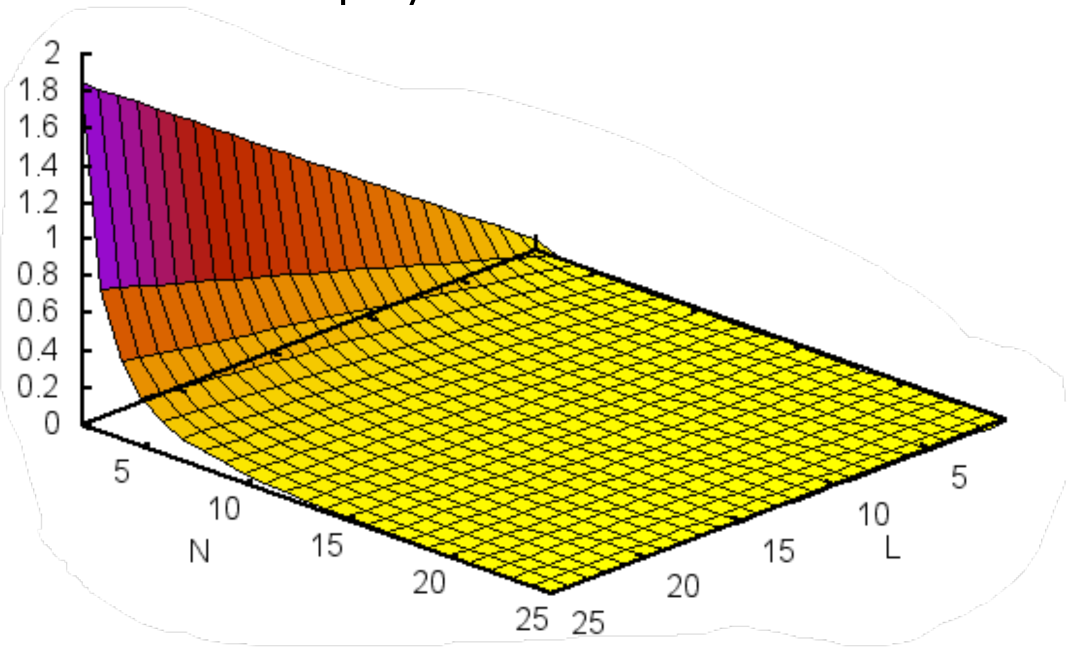

(note: the paper also contains plots for the maximizing (cost of protection, cost of insurance) pairs)

- Payoff difference increases with the potential losses
- Payoff difference decreases when the number of players increases
  - Unless losses are in $L \approx O(N^2)$

Best Shot, Payoff Ratio

Best-shot: Payoff ratio as a function of number of players N and losses L

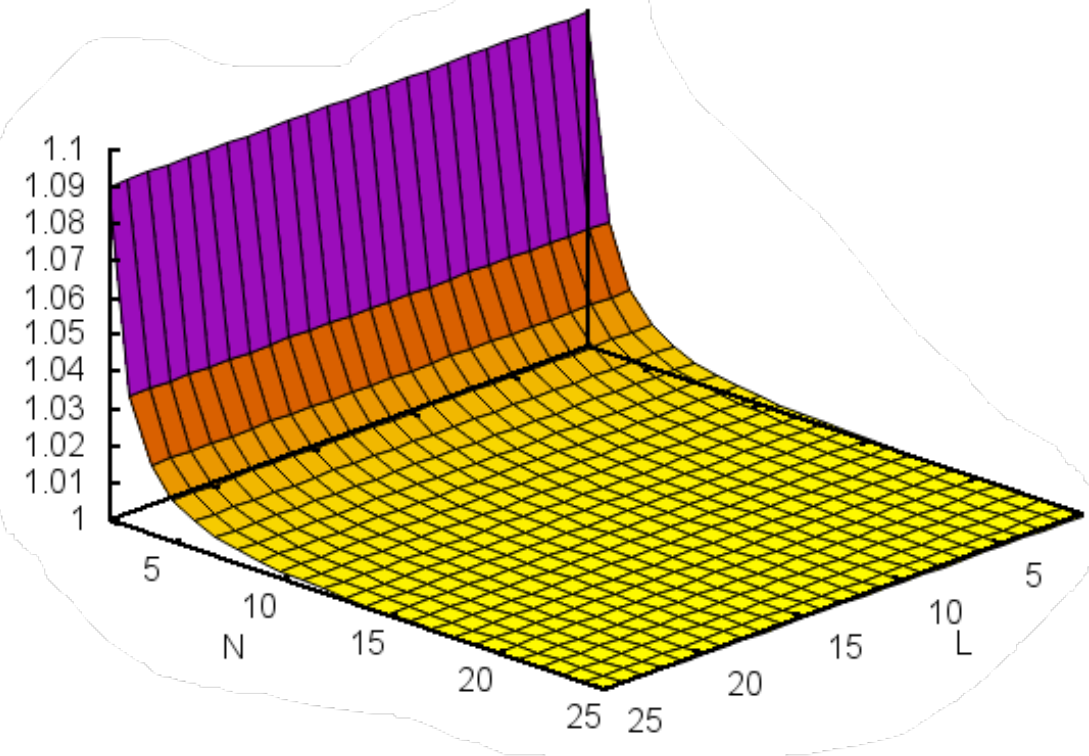

- Payoff ratio independent of L
- Payoff ratio decreases when the number of players increases
- Fairly insignificant overall!
  - 10% at most
- Not shown here: cost ratio metric *always* equal to zero! (significant?!)

Weakest-Link Game

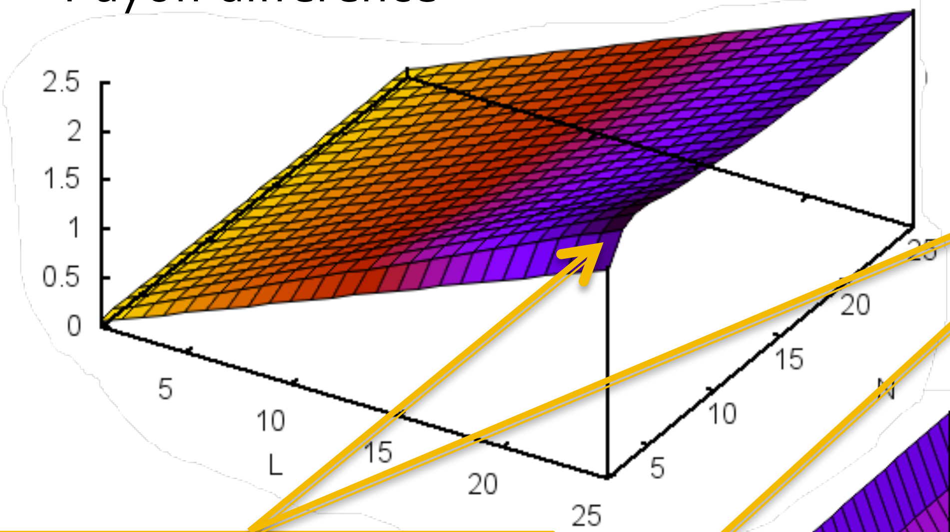

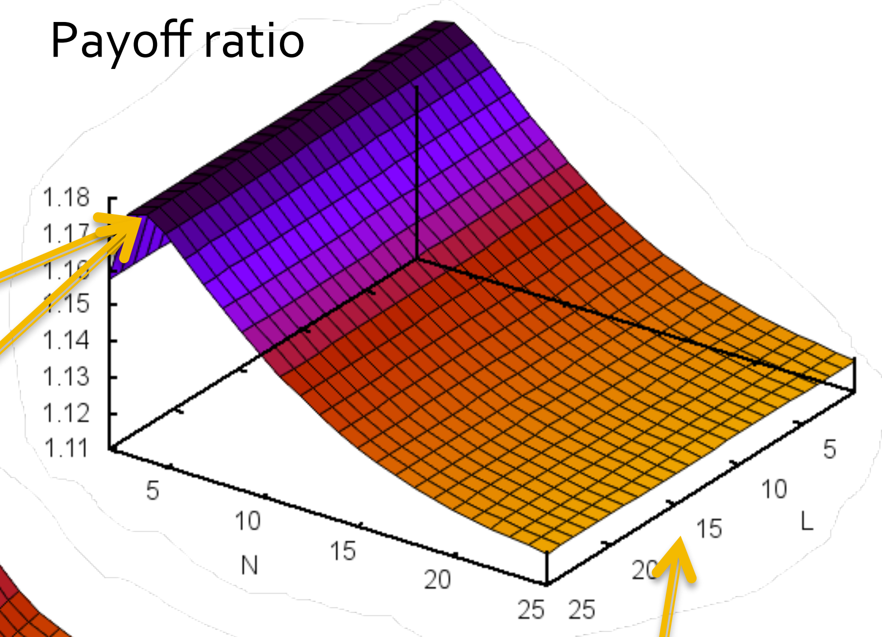

Highest value for 4 players

Slightly more significant (18%), but not catastrophic

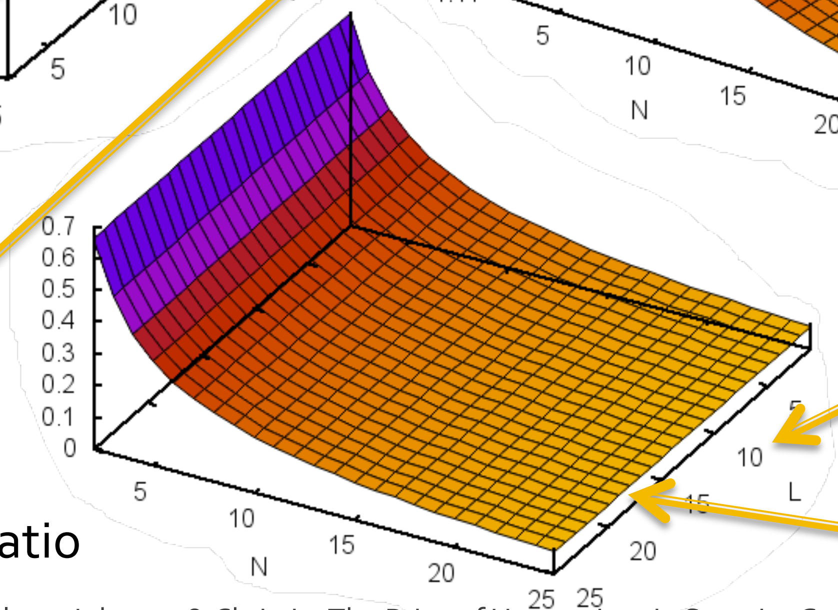

Cost ratio

Independent of L

Goes to zero?!

Total Effort Game

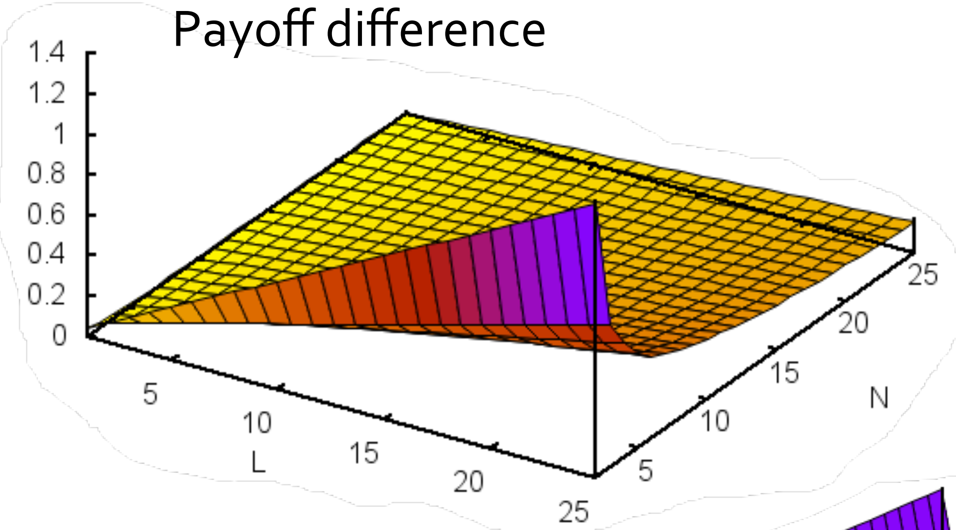

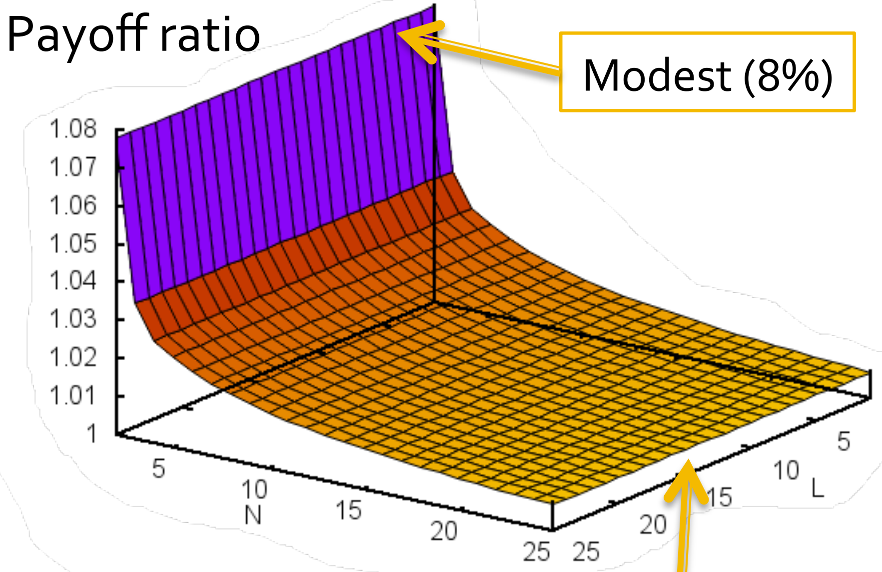

Modest (8%)

Independent of L

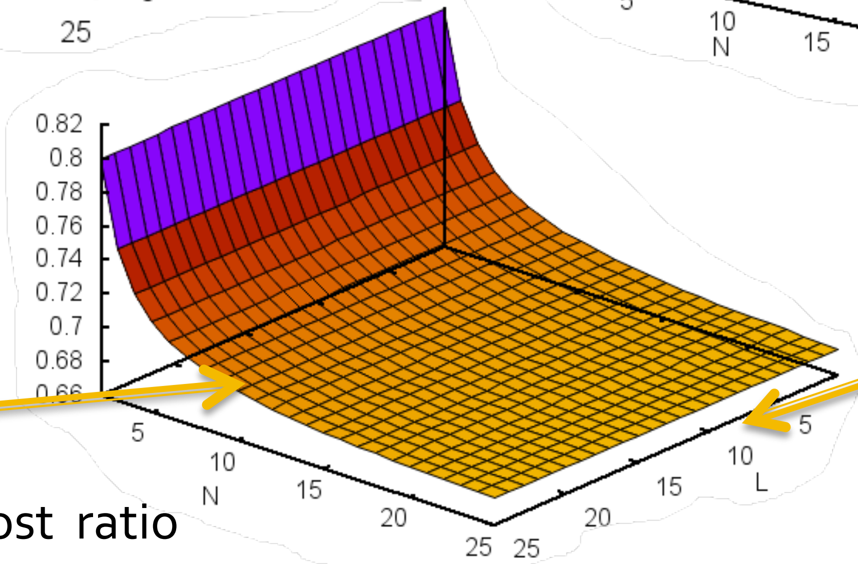

Gets worse in the number of players?!

Finding: Cost Ratio Is Harmful

- Cost ratio metric always inappropriate in all three scenarios
  - Computing ratios of *very* small quantities
  - A penny divided by a dime yields a 0.1… (remember, going to zero is worse)
  - … but is not characteristic of large costs!
    - The fact we are dealing with very small quantities is more important
- Behavioral research has shown robust evidence for consumers' preferences for benefits that are presented as large ratios in comparison to small ratios
  - Useful for marketing snake oil, but not for much else

Finding: Uncertainty vs. Expertise

- All metrics show that uncertainty does not significantly penalizes an expert player
- The **more players** in a network, the **less uncertainty** matters
- **Naïve** strategies have a significantly more disastrous impact on payoffs
  - Not shown today
  - Please see paper and related, companion technical report CMU-CyLab-2009-04

Questions?

The Price of Uncertainty in Security Games

J. Grossklags, B. Johnson and N. Christin
jensg@ischool.berkeley.edu
johnsonb@andrew.cmu.edu
nicolasc@andrew.cmu.edu

Related papers:
http://www.andrew.cmu.edu/user/nicolasc/papers-topic.html

1. Security and Insurance Management in Networks with Heterogeneous Agents [ACM EC'08]
2. Secure or Insure? A Game-Theoretic Analysis of Information Security Games. [WWW'08]
3. Predicted and Observed User Behavior in the Weakest-Link Security Game. [USENIX UPSEC'08]