Decision Analysis for Risk (and Reliability)

Basic concepts and some open problems

David Ríos Insua

david.rios@urjc.es

Royal Academy of Sciences, Spain

Belgirate, May ‘11

Goals

- Provide a review of key concepts and issues in decision analysis
- Illustrate their relevance in risk analysis (and reliability analysis)
- Combine with game theory to talk about adversarial risk analysis
- Present some open problems

Outline

- Review of decision analysis
- Decision analytic framework for risk analysis.
- Aviation case study
- Adversarial risk analysis
- Some open problems

Decision Analysis

- Purpose

Support a Decision Maker in making a decision under uncertainty (consequences of decisions not known with certainty)

Risk Analysis. What is the best security resource allocation in a city?

Reliability Analysis. What is the best SW/HW maintenance policy for the ERP of a university?

DA and Risk Analysis

Risk Analysis. What is the best security resource allocation in a city?

City as a map with cells

For each cell a value

For each cell a forecast of delictive acts

Allocate security resources (constraints)

For each cell forecast impact of resource allocation

Optimal resource allocation

NB: The bad guys also operate intelligent and organisedly!!!

DA and Reliabiltiy Analysis

Reliability Analysis. What is the best SW/HW maintenance policy for the ERP of a university?

Model the HW/SW system (HW and SW blocks interacting)

Reliability forecast for blocks

Reliability forecast for system

Device maintenance policies

Forecast their impact on reliability (and consequently costs)

Optimal maintenance policy

NB: Again, what about bad guys attacking our system??

Decision analysis (cycle)

- Structure problem: Identify alternatives, uncertainties and consequences
- Elicit probabilities, Possibly update in light of data
- Elicit utilities
- Compute alternative with maximum (posterior) expected utility
- Perform sensitivity analysis
- Possibly iterate, until implementation

$$\max_a \int u(c(a,\theta))p(\theta|x)d\theta$$

Modelling preferences. Risk attitudes

- Under certain conditions (Von Neumann-Morgenstern axioms) model preferences and risk attitudes of a person with a utility function.

- Maximum expected utility principle

$$\max_a \int u(c(a,\theta))p(\theta|x)d\theta$$

Modelling preferences. Risk attitudes

- Three basic risk attitudes

  - Risk aversion. Concave utility
  - Risk proneness. Convex utility
  - Risk neutrality. Linear utility

Consequence varying risk attitudes

Risk attitudes

Compare lotteries A and B. Which one do you prefer?

$$\mathbf{A} = \begin{pmatrix} 1/2 & 1/2 \\ 0 & 100000 \end{pmatrix} \text{ y } \mathbf{B} = \begin{pmatrix} 1 \\ 50000 \end{pmatrix}.$$

Risk attitudes

Compare lotteries A and B. Which one do you prefer?

$$\mathbf{A} = \begin{pmatrix} 1/2 & 1/2 \\ 0 & 100000 \end{pmatrix} \text{ y } \mathbf{B} = \begin{pmatrix} 1 \\ 50000 \end{pmatrix}.$$

If you prefer B to A, you seem to prefer the sure prize to the lottery (risk aversion)

In such case, the expected utility of B should be bigger than that of A:

$\frac{1}{2} * u(0) + \frac{1}{2} * u(100000) = \frac{1}{2} < 1 * u(50000)$

Graphically...

Risk attitudes

Compare lotteries A and B. Which one do you prefer?

$$\mathbf{A} = \begin{pmatrix} 1/2 & 1/2 \\ 0 & 100000 \end{pmatrix} \text{ y } \mathbf{B} = \begin{pmatrix} 1 \\ 50000 \end{pmatrix}.$$

If you prefer A to B, you seem to prefer the lottery to the sure prize (risk proneness)

In such case, the expected utility of A should be bigger than that of B:

$\frac{1}{2} * u(0) + \frac{1}{2} * u(100000) = \frac{1}{2} > 1 * u(50000)$

Graphically...

Risk attitudes

Compare lotteries A and B. Which one do you prefer?

$$\mathrm{A} = \begin{pmatrix} 1/2 & 1/2 \\ 0 & 100000 \end{pmatrix} \text{ y } \mathrm{B} = \begin{pmatrix} 1 \\ 50000 \end{pmatrix}.$$

If you are indifferent between A and B, you seem indifferent between the lottery and the sure prize (risk neutrality)

In such case, the expected utilities of A and B should be equal:

$\frac{1}{2} * u(0) + \frac{1}{2} * u(100000) = \frac{1}{2} = 1 * u(50000)$

Graphically...

Which risk attitude??

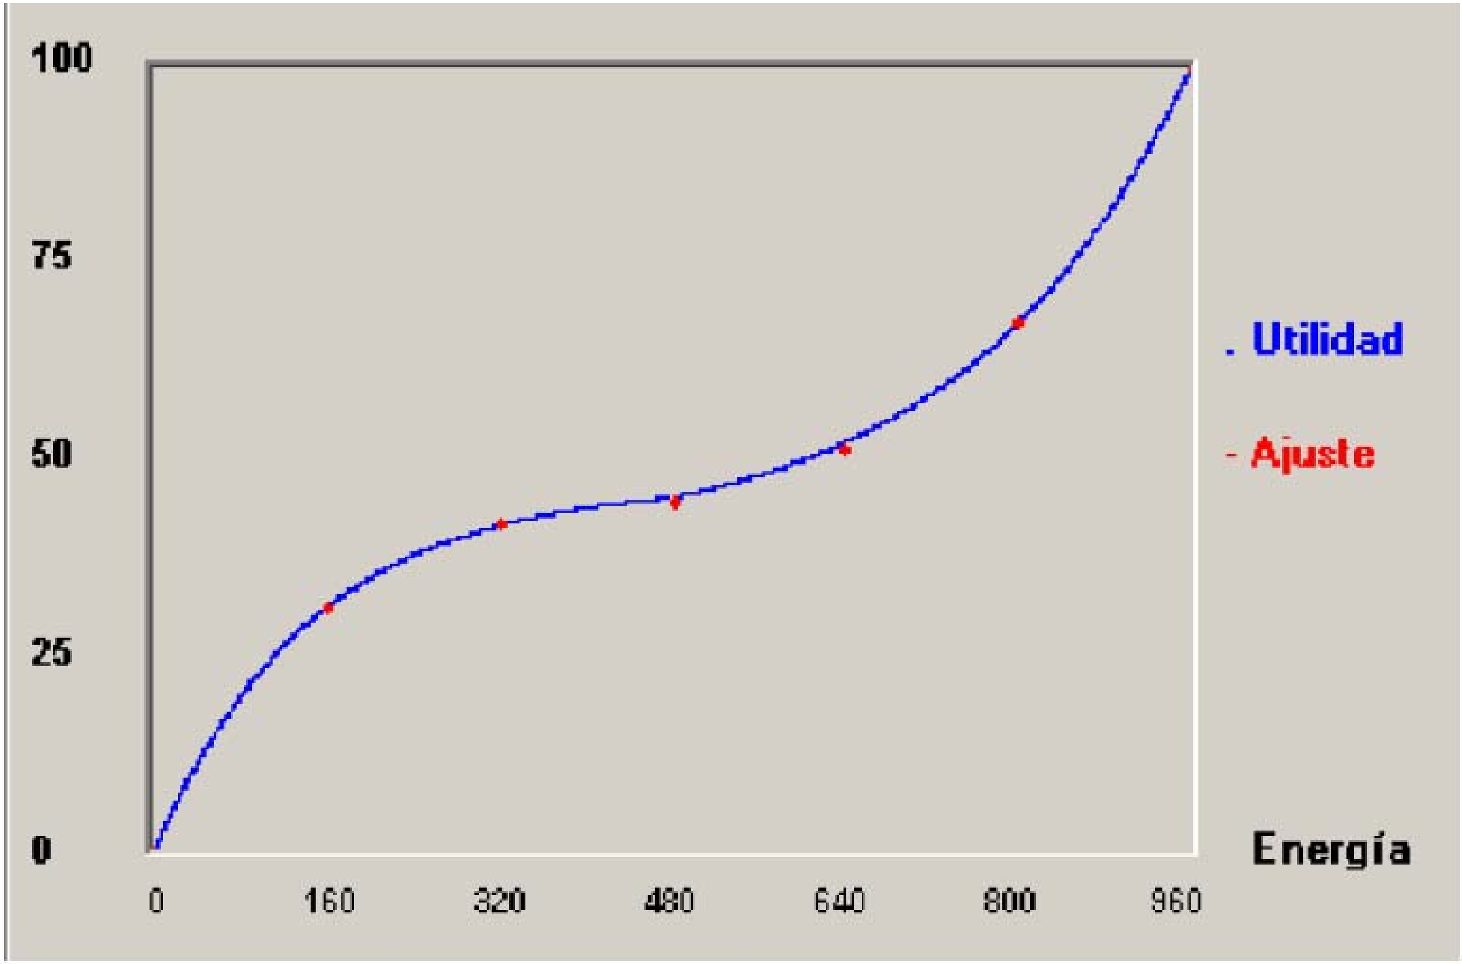

Influence Diagrams

- Tool to structure (and solve) decision making problems
- Direct acyclic graph G=(N,A)
- Three main types of nodes.
  - Chance. Circle
  - Decision. Square
  - Value. Hexagon, Diamond
    - Fourth type of node. Deterministic. Double circle
- Two types of arcs
  - Arcs into decision nodes
  - Arcs into chance and value nodes

Influence Diagrams

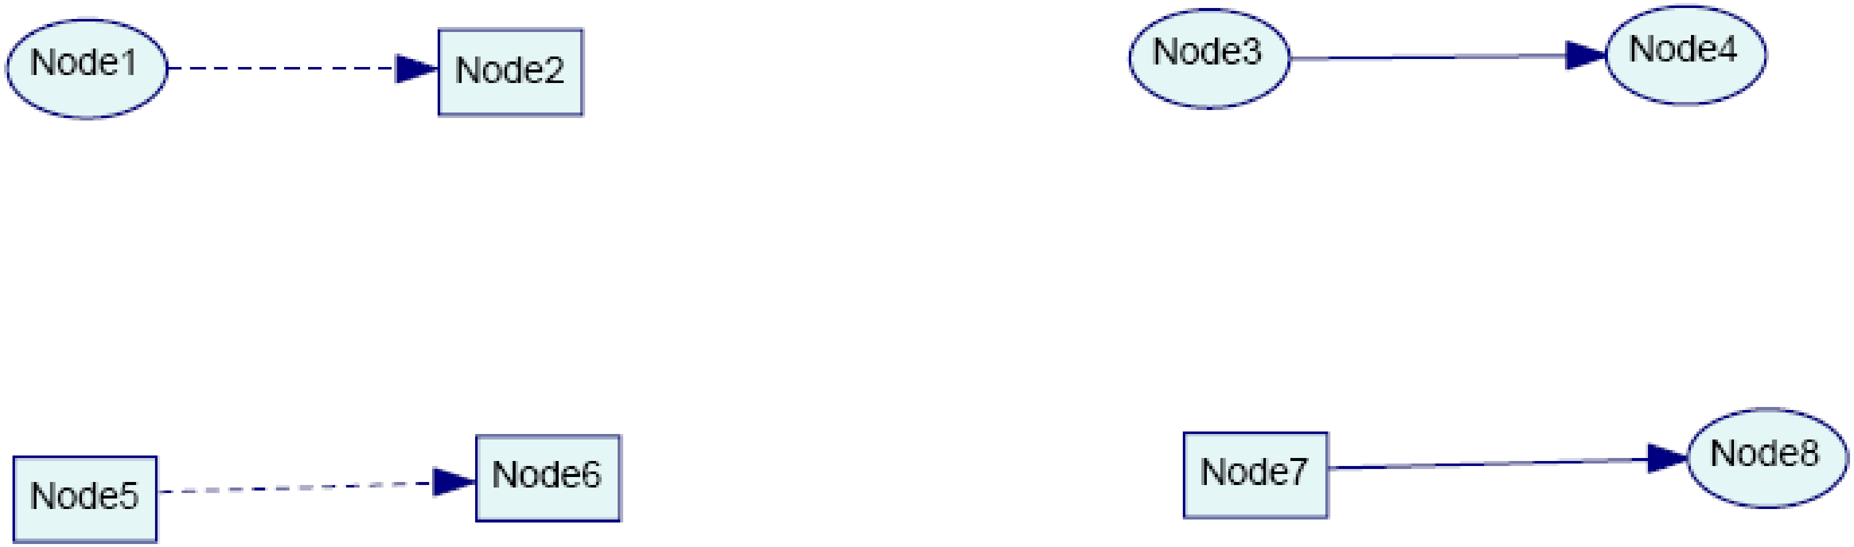

Influence diagrams: Example

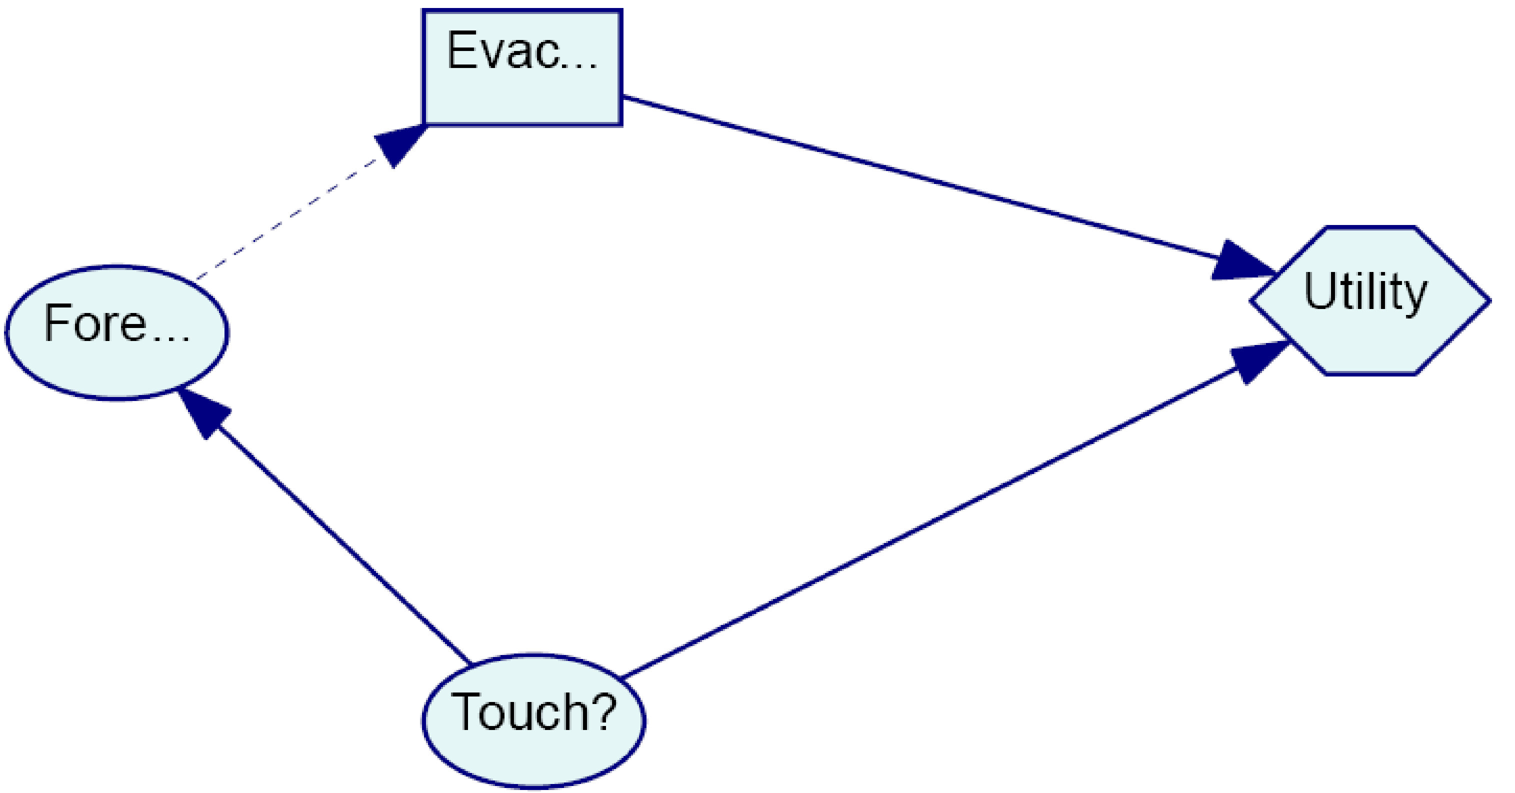

Influence diagrams: Multiple objectives

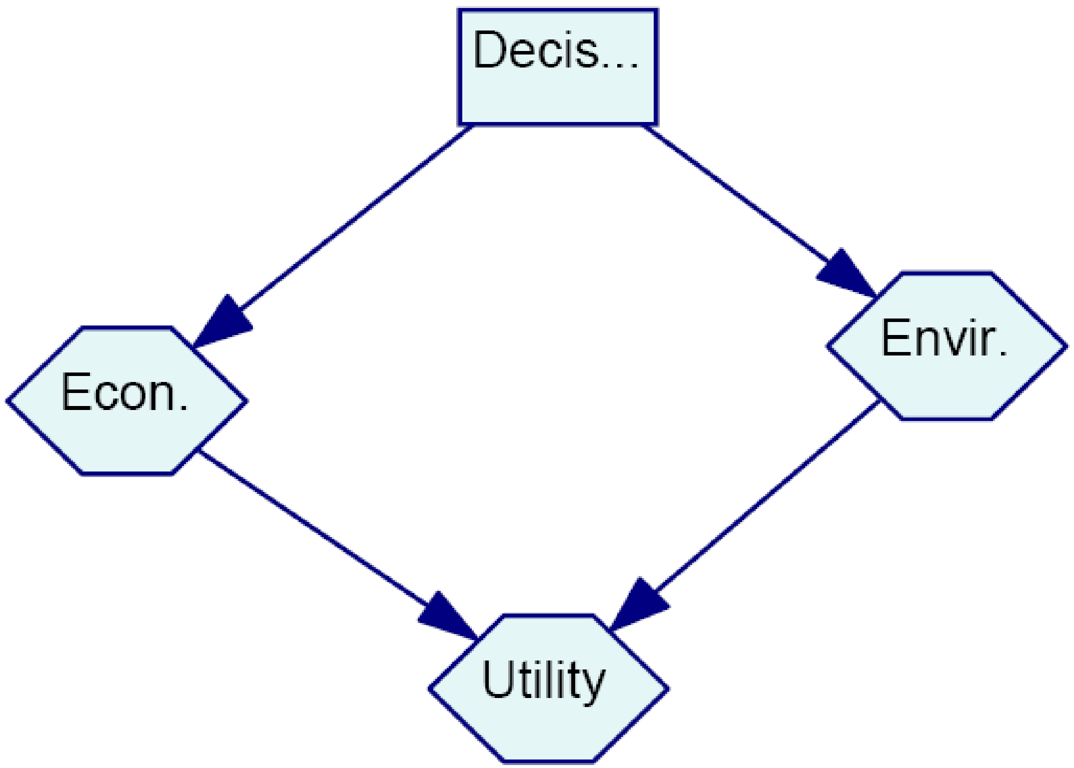

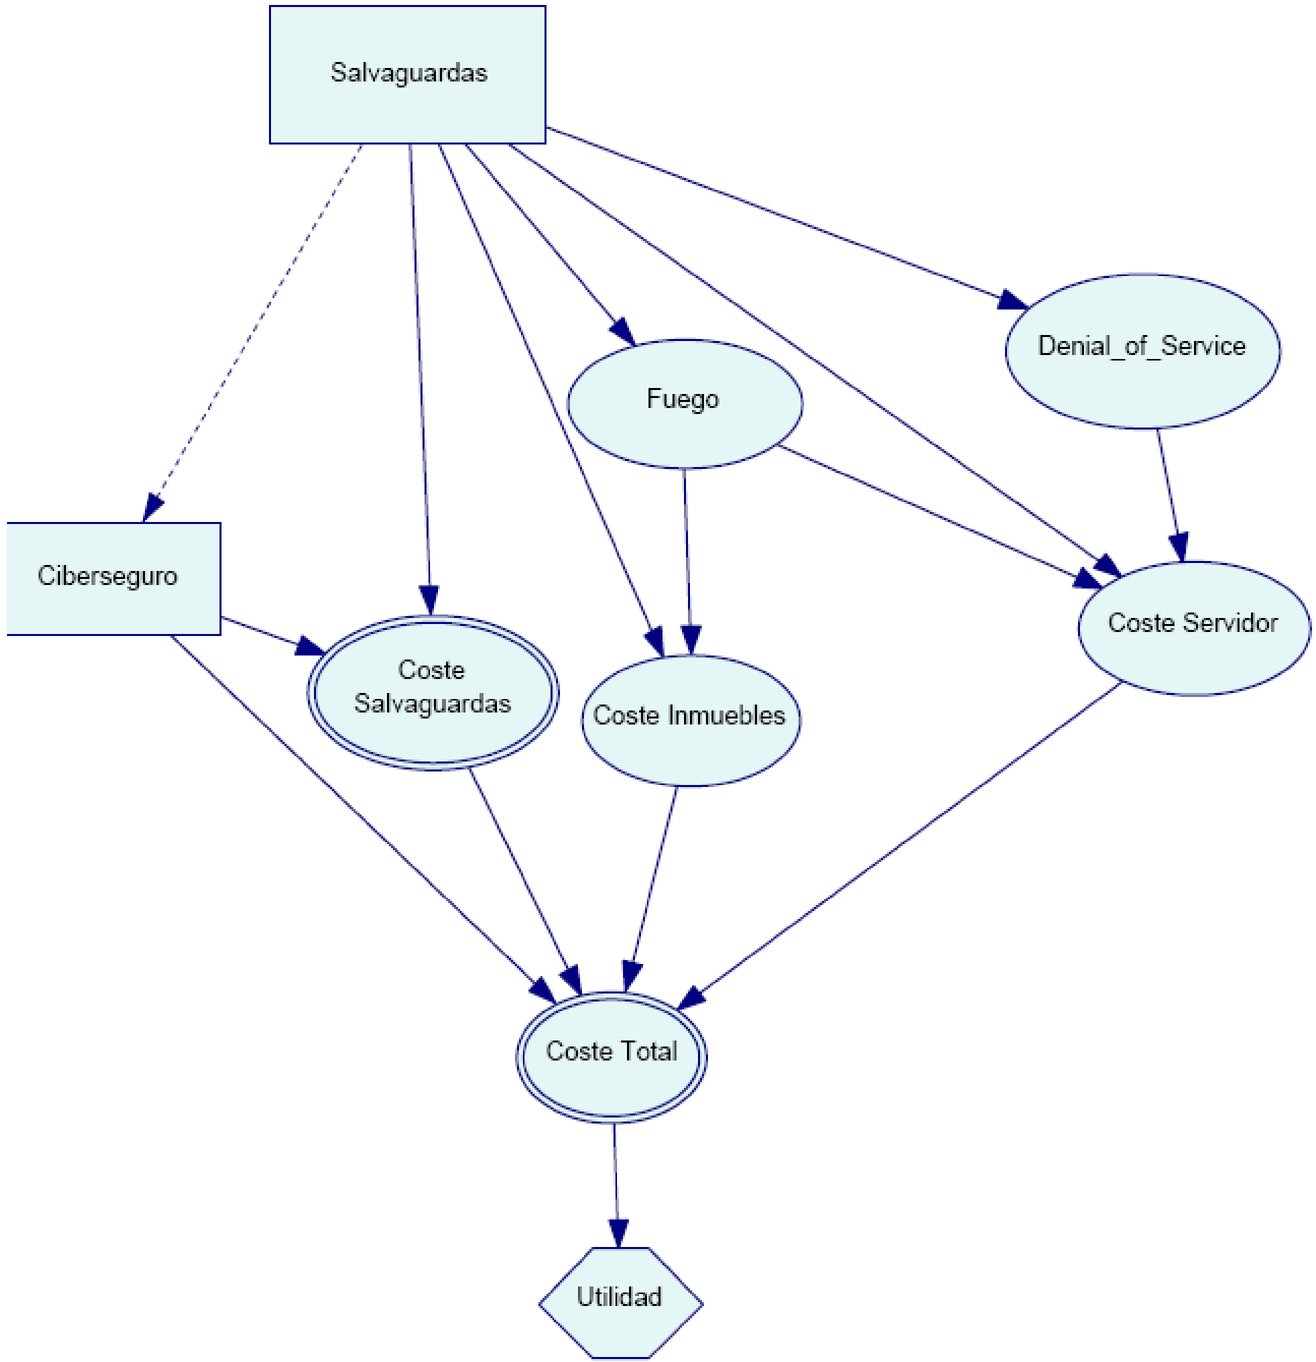
Salvaguardas
Denial_of_Service
Fuego
Ciberseguro
Coste Servidor
Coste Salvaguardas
Coste Inmuebles
Coste Total
Utilidad

Bayesian computational methods

In general, we shall need to compute (posterior) maximum expected utility alternatives

$$\max_a \int u(c(a,\theta))p(\theta|x)d\theta$$

Sometimes, it may be convenient to solve

$$\max_a \int u(a,\theta)p(x|\theta)p(\theta)d\theta$$

One possibility, approximate expected utilities by Monte Carlo then optimise the MC sums... Sampling from the posterior??

1. Select a sample $\theta^1, ..., \theta^m \sim p(\theta|x)$.

2. Solve the optimisation problem

$$\max_{a \in \mathcal{A}} \frac{1}{m} \sum_{i=1}^{m} u(a, \theta^i)$$

yielding $a_m(\theta)$.

Computational methods: Gibbs sampler

In some contexts, we are not able to sample directly from the target distribution, but we may sample from the marginal conditionals. Then, under appropriate conditions, the following scheme is designed to converge to the target distribution

1. Choose initial values $(\theta_2^0, \ldots, \theta_k^0)$. $i = 1$

2. Until convergence is detected, iterate through

   . Generate $\theta_1^i \sim \theta_1 | \theta_2^{i-1}, .., \theta_k^{i-1}$

   . Generate $\theta_2^i \sim \theta_2 | \theta_1^i, \theta_3^{i-1}, ..., \theta_k^{i-1}$

   . ...

   . Generate $\theta_k^i \sim \theta_k | \theta_1^i, ..., \theta_{k-1}^i$.

   . $i = i + 1$

Computational methods: Gibbs sampler

Imagine we need to sample from

$$p(\theta_1, \theta_2 \mid x) = \frac{1}{\pi} \exp\{-\theta_1(1+\theta_2^2)\}$$

The conditionals are easily identified and a Gibbs sampler scheme is

1. Choose initial value for θ_2; e.g., the posterior mode, $\theta_2^0 = 0$. $i = 1$

2. Until convergence, iterate through
   - Generate $\theta_1^i = \mathcal{E}/(1+[\theta_2^{i-1}]^2)$, ($\mathcal{E}$, standard exponential).
   - Generate $\theta_2^i = Z/\sqrt{2\theta_1^i}$, ($Z$, standard normal).

Computational methods: Metropolis sampler

Sometimes, we cannot sample from the conditionals. However, as we know, up to a constant, the target distribution, by choosing an appropriate candidate generating distribution q(.|.), under appropriate conditions, the following scheme is designed to converge to the target distribution

1. Choose initial values θ^0. $i = 0$
2. Until convergence is detected, iterate through
   - Generate a candidate $\theta^* \sim q(\theta|\theta^i)$.
   - If $p(\theta^i)q(\theta^i \mid \theta^*) > 0$, $\alpha(\theta^i, \theta^*) = \min\left(\frac{p(\theta^*)q(\theta^*|\theta^i)}{p(\theta^i)q(\theta^i|\theta^*)}, 1\right)$;
   - else, $\alpha(\theta^i, \theta^*) = 1$.
   - Do

$$\theta^{i+1} = \begin{cases} \theta^* & \text{with prob } \alpha(\theta^i, \theta^*), \\ \theta^i & \text{with prob } 1 - \alpha(\theta^i, \theta^*) \end{cases}$$

   - $i = i + 1$.

Computational methods: Augmented probability simulation

Frequently, the involved posterior depends on decision made. The following observation helps in this context. Define an artificial distribution such that (u, needs to be nonnegative)

$$h(a,\theta) \propto u(a,\theta) \times p_\theta(\theta \mid x, a).$$

Then, the marginal of the artificial is proportional to expected utility

$$h(a) = \int h(a,\theta)\, d\theta_{a,\theta} \propto \Psi(a).$$

This suggests the scheme

1. Generate a sample $((\theta^1, a^1), ..., (\theta^m, a^m))$ from density $h(a, \theta)$.
2. Convert it to a sample $(a^1, ..., a^m)$ from the marginal $h(a)$.
3. Find the sample mode.

Outline

- Review of decision analysis
- Decision analytic framework for risk analysis.
- Aviation case study
- Adversarial risk analysis
- Some open problems

Risks in modern world

Risk management is the top prority for top management in major companies. Four years ago: human resources and talent management (Accenture report)

Demands for security in an increasingly globalised economy, pressure of regulators,

Risk analysis: What for??

Risk management for an existing or proposed facility

Development of regulations

Demonstration of compliance with regulations

Demonstration of need for further improvement

Litigation

Scientific enquiry

The risk management process

1. Determination of objectives

   Preserve the operating effectiveness of the organisation

2. Identification of risks
3. Evaluation of risks
4. Considering alternatives and selecting the risk treatment device
5. Implementing the decision
6. Evaluation and review

A framework for risk analysis: starting assumptions

- Only interested in costs…
- An existing alternative
- Just my organisation is relevant
- Aim. Maximise expected utility

Risk analysis framework

- Forecast **costs** under normal circumstances
- Identify hazard events, estimate probabilities and impacts on costs (additional induced costs)
- Forecast costs (**a “mixture” model**). Compute changes in expected utility. If too big,…
- Identify interventions, estimate impact on probabilities and/or costs.
- Compute expected utilities. Choose best intervention (if gain is sufficient)

Basic setting

- Design given (no interventions, status quo)
- (Random) costs are identified
- Expected utility computed

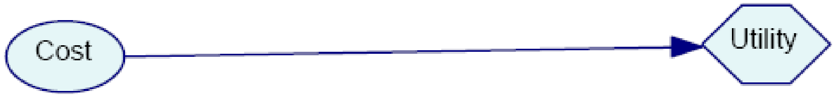

$$\Psi = \int u(c)\pi(c)dc$$

Basic setting

- Design given

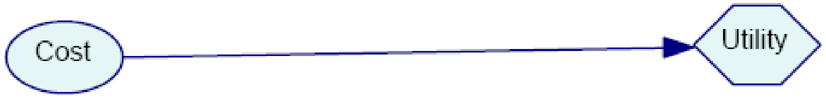

$$\Psi = \int u(c)\pi(c)dc$$

- Including design choice

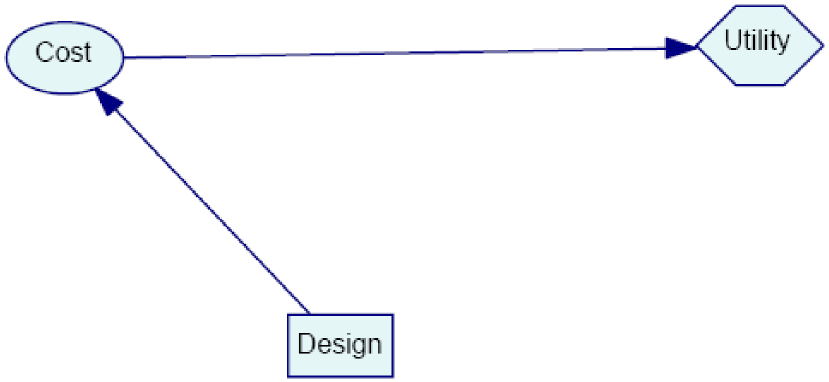

$$\max_a \Psi(a) = \int u(c)\pi(c|a)dc$$

Risk assessement

- Likelihood and impact of identified hazards. They happen with a certain probability and entail an additional cost

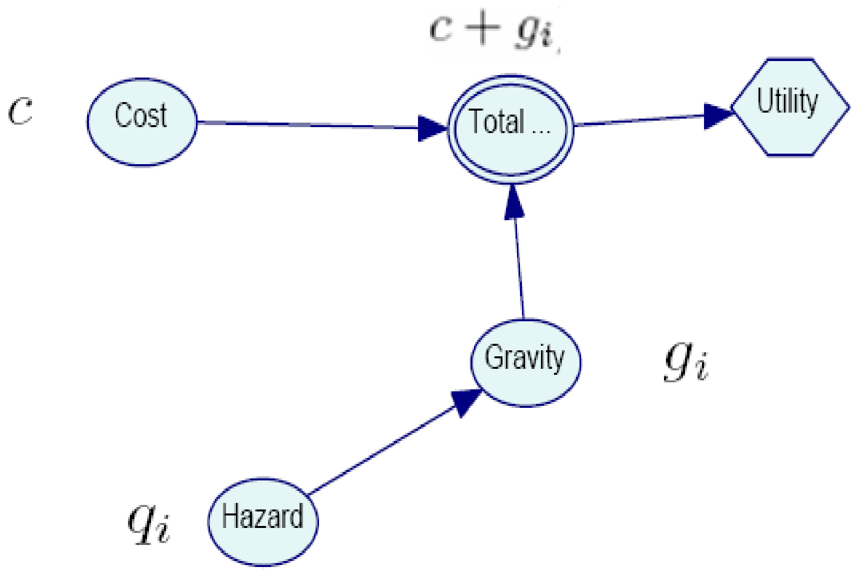

- Compute expected utility after risk assessed:

$$\Psi_r = \int \int \int \sum q_i u(c + g_i) \pi(q) \pi(g) dq dg \ \pi(c) dc$$

- Impact of risks: $\Psi - \Psi_r$

If impact is too high, we need to manage risks

Risk management

- Intervention to be chosen:

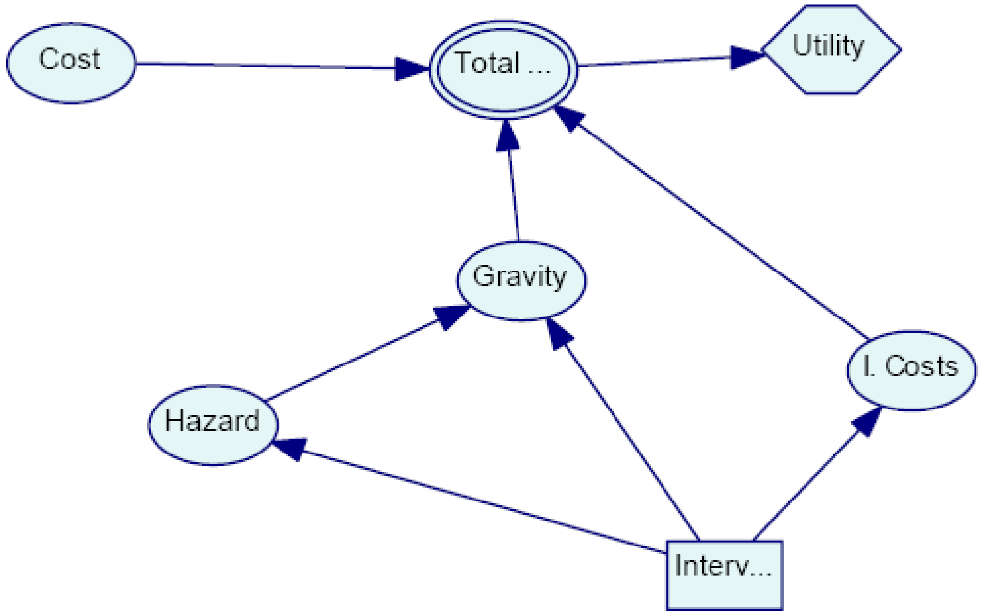

Interventions tend to reduce the likelihood of hazard appearance and its gravity... but they also entail a cost

$$\Psi_d = \max_d \Psi_r(d) = \max_d \int\int\int\int \sum q_i u(c+g_i+c_d)\pi(q|d)\pi(g|d)dq\, dg\pi(c)\pi(c_d) \;\; dc_d\, dc$$

- Gain through managed risk: $\Psi_d - \Psi_r$

Choose the intervention which provides the biggest gain, if it is sufficiently big...

Outline

- Introduction
- Review of decision analysis
- Decision analytic framework for risk analysis.
- Aviation case study
- Adversarial risk analysis
- Some open problems

An example: Unintended slide deployment

An example from aviation operations. Background

- Safety critical in aviation industry
- Increasing competition forcing cost reduction, even more under crisis
- Relatively simple tools for risk analysis commercial aviation operations

- Unintended slide deployment under normal operations within a comercial airline
- Inflatable slides to facilitate passenger evacuation in emergency situations
- (Expected??) cost 20 million USD/year for the whole industry (IATA, 2000)

An example: Unintended slide deployment

- Factors affecting incidents
- **Severity analysis (cost)**
- Risk assessment
- Countermeasures?
- Best countermeasure: risk management

- Estimated annual savings 600000 €

An example: Unintended slide deployment

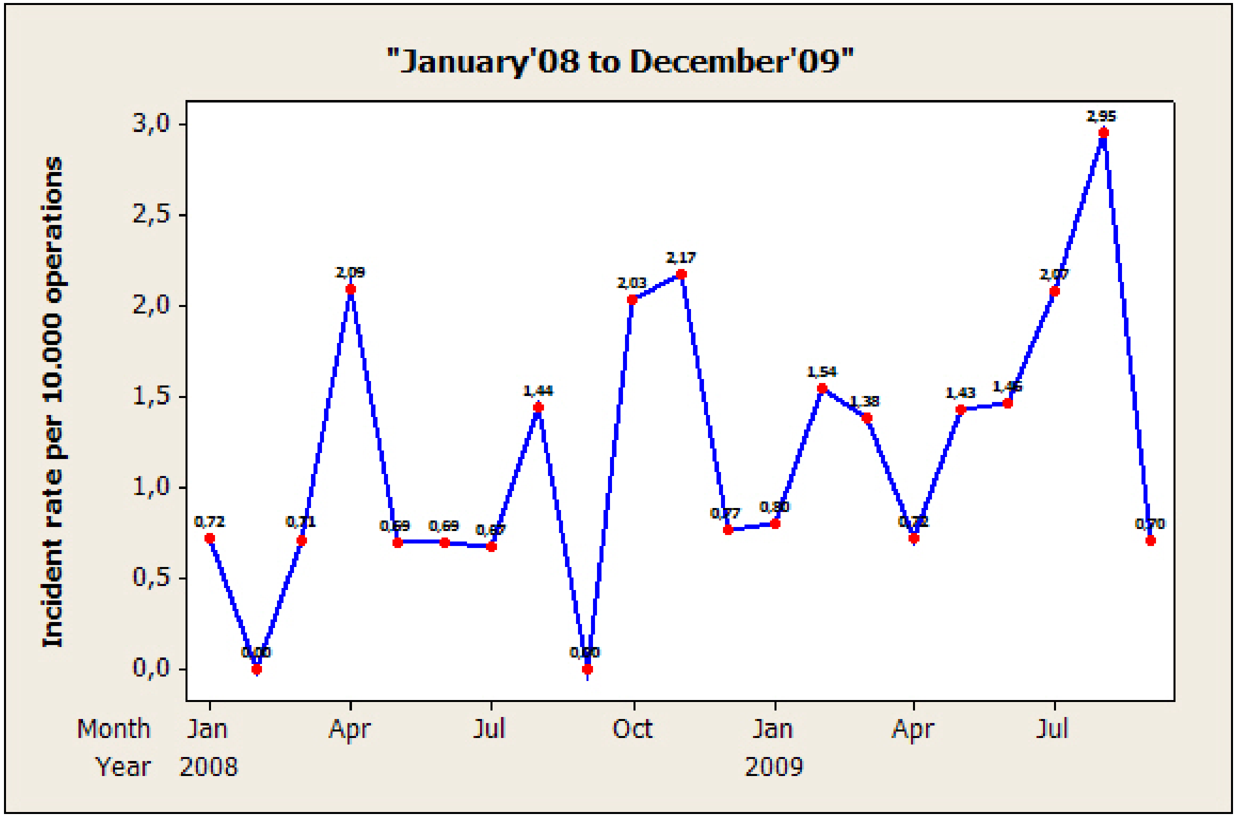

An example from aviation operations. Incident analysis

- The following potentially factors are identified

Factor	Relevance	Factor levels
Aircraft type	Yes, Moderate	A $>$ B
Airport	No	Nearly 50
Pairing day	Yes	First $>$ Second $>$ Third
Flight turn	Yes	First $>$ (Second, Third)

We build a logistic regression model with three explanatory variables

Relevant operational phase and personnel involved

Factor	Relevance ranking
Operational phase	Arrival $>$ Departure $>>$ Refueling $>$ Preflight $=$ Stopover
Staff involved	(A, B) $>$ (C,D,E,F,G,H,I)

Finally, 7 errors, 9 procedure interruptions, 19 procedure non compliances (Dirichlet model)

An example: Unintended slide deployment

- Cost
  - Removal cost
  - Transportation cost
  - Repair cost
  - Ground delay associated costs

An example: Unintended slide deployment

- Removal Cost
  - Lab x T_m
  - $T_{m.}$ Expert assesses min (30), max (60), most likely (45). Adjust triangular distribution with 0.05, 0.95 quantiles at min, max . Tri (0.385,0.75,1.115)
- Transportation cost

An example: Unintended slide deployment

- Maintenance cost

$$C_m = q \times C_m^i + (1-q) \times C_m^e$$

- q assessed Beta (16,4)
- C_m

	Bf	Ba	Bw	B2/3
Int. costs	1840	1480	1630	1430
Ext. costs	2605	2323	4571	4741

	A1	A2	A3	A6	A6w
Int. costs	4160	4040	2400	3630	3210
Ext. costs	6429	4850	5785	7423	4946

	Bf	Ba	Bw	B2/3
Incidents	17	4	1	5
Parameters	18	5	2	6

	A1	A2	A3	A6	A6w
Incidents	4	2	1	0	0
Parameters	5	3	2	1	1

An example: Unintended slide deployment

- Costs in relation with delays

$$T_d = p_0\ I_0 + p_1\ F_d$$

$$p_0 + p_1 = 1$$

$$p_0, p_1 \geq 0$$

$$p_0 | data \sim Be(14, 23)$$

$$F_{d_B} \sim Wei(\theta = 0, \alpha, \beta)$$

$$F_{d_A} \sim p\ Wei(\theta = 0, \alpha, \beta) + (1-p)\ Wei(\theta, \alpha, \beta),$$

$$f(x \mid \theta, \alpha, \beta) = \alpha \frac{(x-\theta)^{\alpha-1}}{\beta^{\alpha}}\ \exp(-((x-\theta)/\beta)^{\alpha})$$

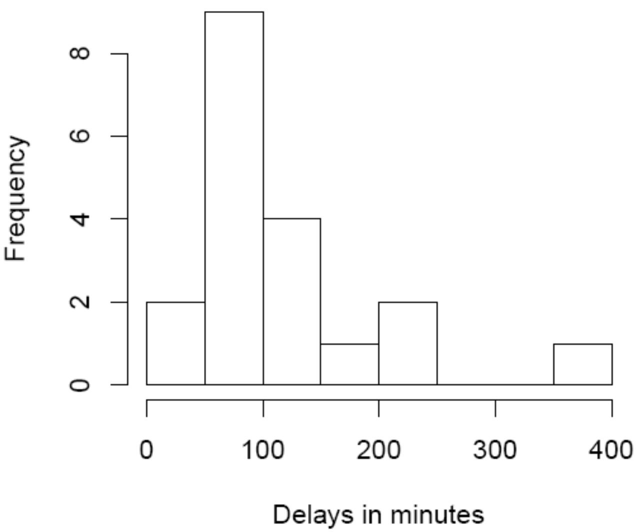

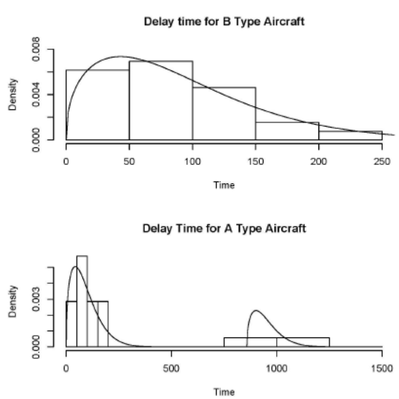

An example: Unintended slide deployment

- Costs in relation with delays

	A Flights	B Flights
	(Min, most likely, max)	(Min, most likely, max)
Passenger Hard Costs	(0.12, 0.19, 0.24)	(0.12, 0.19, 0.24)
Passenger Soft Costs	(0.06, 0.19, 0.22)	(0.06, 0.19, 0.22)
Marginal Crew Costs	(0.00, 14.00, 39.00)	(0.00, 7.90, 16,59)
Marginal Maintenance Costs	(0.65, 0.81, 0.97)	(0.38, 0.47, 0.56)
Total Costs	(0.83, 15.19, 40.27)	(0.56, 8.75, 17.61)

An example: Unintended slide deployment

- Annual costs due to incidents

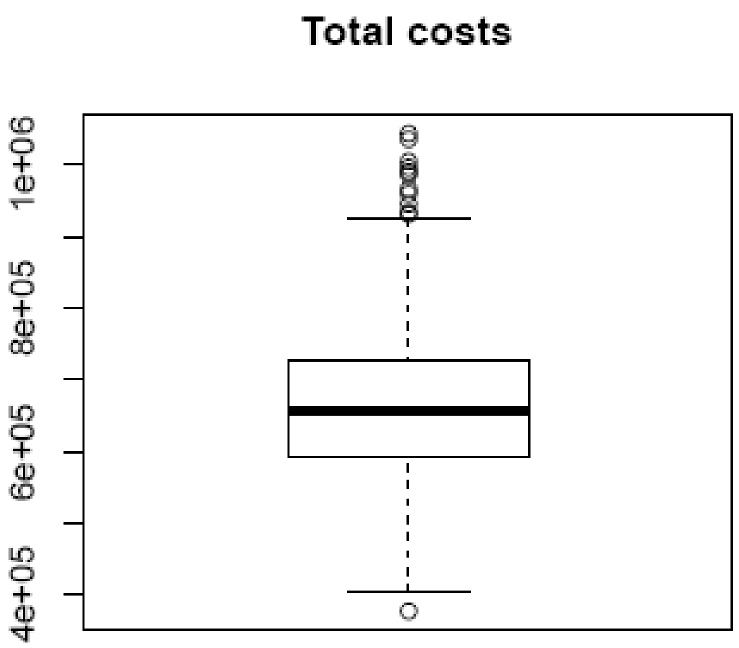

An example from aviation operations: Risk management

- Countermeasures
  - Change procedure (to ‘eliminate’ interruptions and mitigate errors, pratically no cost)
  - Training course to key personnel (to mitigate errors and noncompliances, practically no cost)
  - Awareness campaign to key personnel through newsletters, etc... (same objective, cost 6000 euros)
  - Ligt and sound warning device at each door (to mitigate errors, interruptions andnoncompliances, cost 2500 euros per door) (or only Bf)
  - Visual reminders at each door (to mitigate errors adn noncompliances, cost 120 euros per door) (or only Bf)
- Note that, essentially, we only affect incident likelihood, but not incident severity

An example from aviation operations: Risk management

Countermeasure	1 year	5 years
Procedure revision	252902	1214935
Awareness campaign	524477	2492943
Warning devices, St. 1	1307393	1335514
Warning devices, St. 2	616058	2137866
Visual reminders, St. 1	631403	2881078
Visual reminders, St. 2	677329	3228759
None	663400	1490047

Countermeasure	1 year	5 years
Awareness campaign	123724	567739
Warning devices,St. 1	1302529	1312149
Warning devices, St. 2	352862	873480
Visual reminders, St. 1	273448	1161478
Visual reminders, St. 2	236060	1108918
None	252902	1214935

Outline

- Review of decision analysis
- Decision analytic framework for risk analysis.
- Aviation case study
- Adversarial risk analysis
- Some open problems

Adversarial risks

- Other intelligent participants

  Auctions for large projects, Counterterrorism, Regulators,…

- Their actions influence my risks

- My actions influence their risks

- Some nodes might be shared…

- Possibly conflicting interests, but possibly cooperating,…

Adversarial risks: Me and other

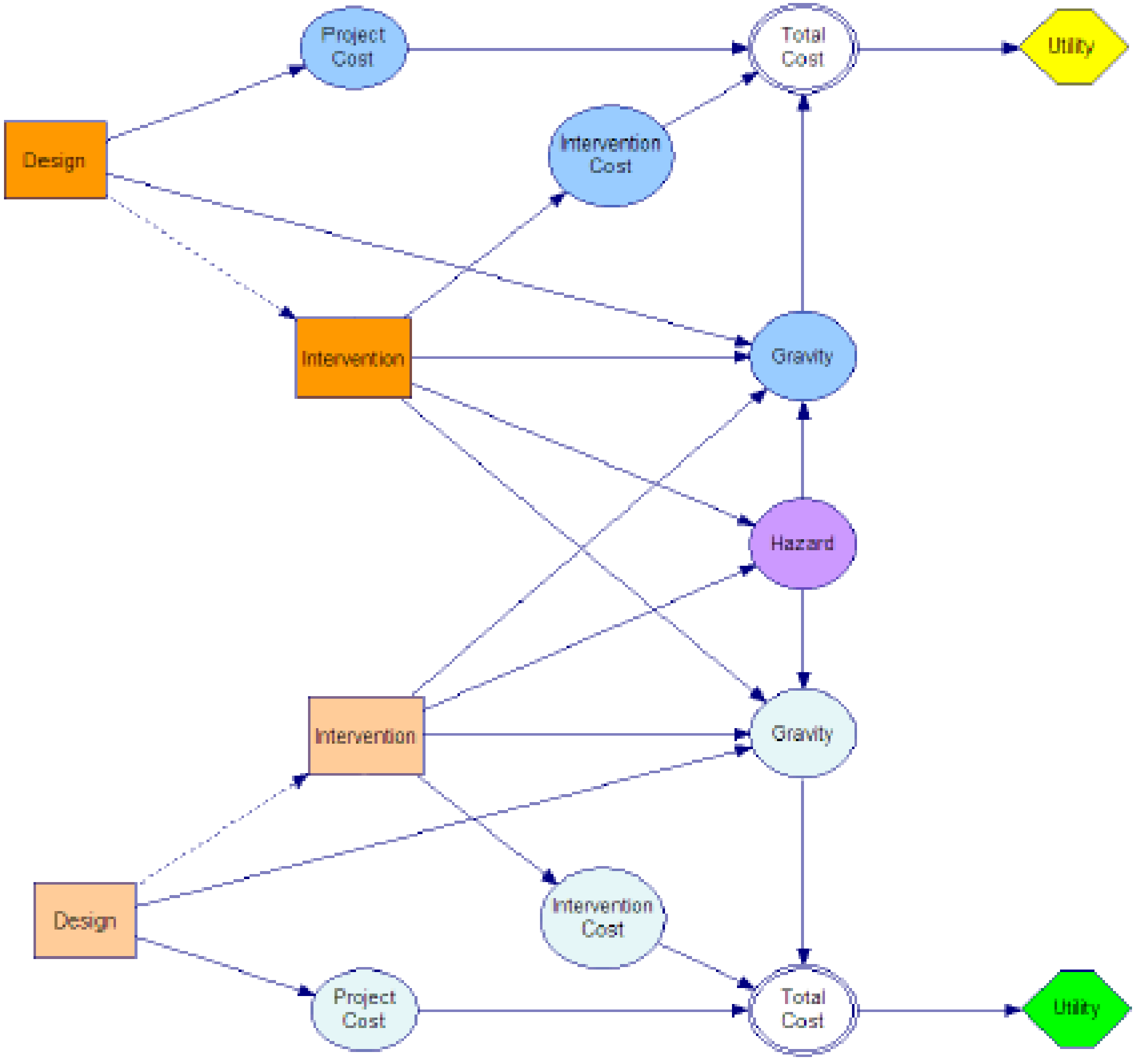

Adversarial risks: Modelling

$$\Psi_1((a_1,d_1)|(a_2,d_2)) = \int\int\int\int \sum q_{1i}^{a_1} u_1(c_1 + g_{1i}^{a_1} + c_{d1}) \pi(q_1^{a_1}|d_1,d_2)\pi(g_1^{a_1}|d_1,d_2) dq_1^{a_1} dg_1^{a_1} \ \pi(c_1|a_1)\pi(c_{d_1}) dc_{d_1} dc$$

$$\Psi_2((a_2,d_2)|(a_1,d_1)) = \int\int\int\int \sum q_{2i}^{a_2} u_2(c_2 + g_{2i}^{a_2} + c_{d2}) \pi(q_2^{a_2}|d_1,d_2)\pi(g_2^{a_2}|d_1,d_2) dq_2^{a_2} dg_2^{a_2} \ \pi(c_2|a_2)\pi(c_{d_2}) dc_{d_2} dc$$

	(a_2, d_2)
(a_1, d_1)	$\Psi_1((a_1,d_1)\|(a_2,d_2)), \Psi_2((a_2,d_2)\|(a_1,d_1))$

One-sided prescriptive support

- Asymmetric prescriptive/descriptive approach (Raiffa)
  - Prescriptive advice to one party conditional on a (probalistic) description of how others will behave
- A Bayesian approach (Kadane, Larkey…)
  - Use a SEU model for supporting the Defender
  - Treat the Attacker's decision as uncertain
  - Help the Defender to assess probabilities of Attacker's decisions
- Adversarial Risk Analysis
  - Weaken common (prior) knowledge assumption
  - Develop methods for the analysis of the adversaries' thinking to anticipate their actions.
    - We assume that the Attacker is a *expected utility maximizer*
    - But other models may be possible

Assessing adversary's intelligent decisions

- Distinction between uncertainty stemming from
  - Nature
  - Intelligent adversaries' actions
- How to assess probabilities of Attacker's actions??
- Assuming the Attacker is a SEU maximizer
  - Based on an analysis of his decision problem
    - Assess Attacker' probabilities and utilities
    - Find his action of maximum expected utility
- Uncertainty about Attacker' decision should reflect
  - Defender's uncertainty about Attacker's probabilities and utilities
- Sources of information
  - Available past statistical data of Attacker's decision behavior
  - Expert knowledge
  - Non-informative (or reference) distributions

Adversarial risks: Bayesian approach

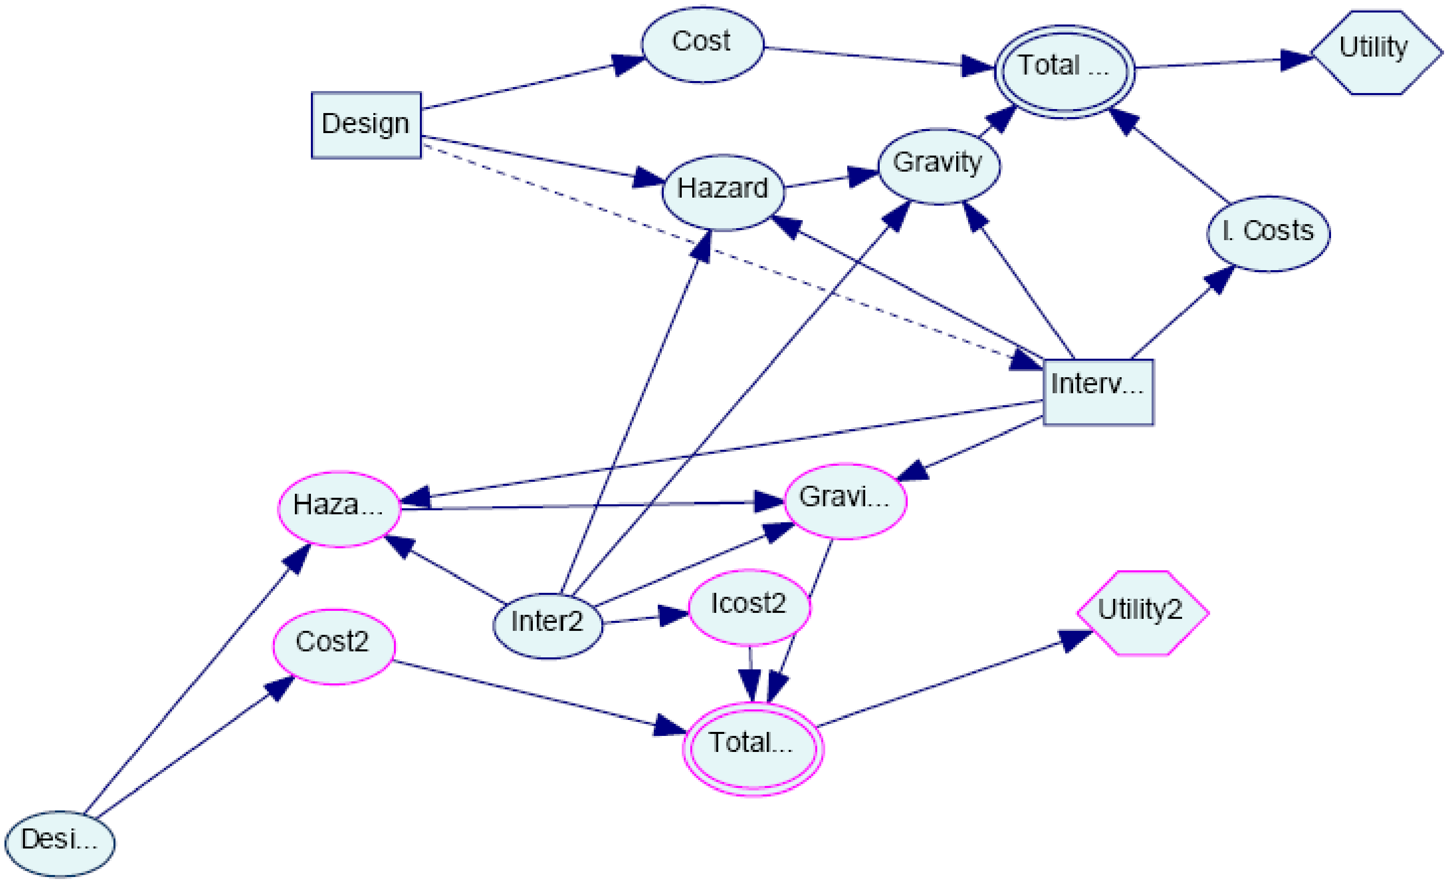

Defend-Attack sequential model

- Two intelligent players
  - Defender and Attacker
- Sequential moves
  - First Defender, afterwards Attacker knowing Defender's decision

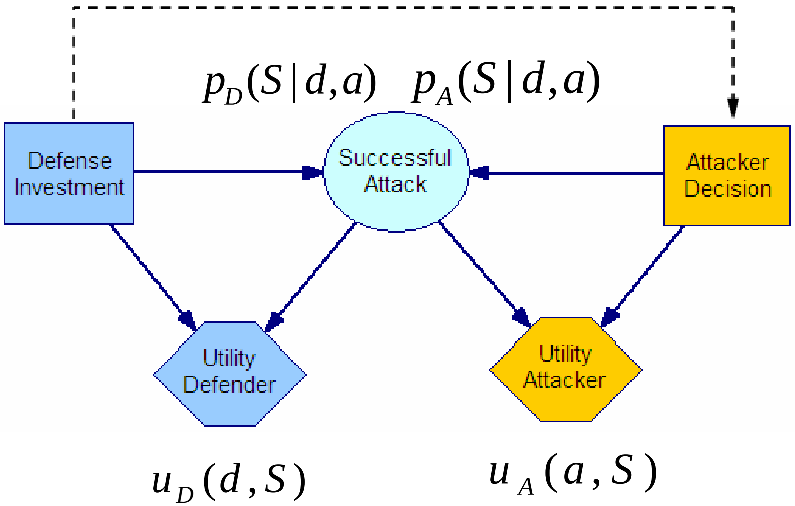

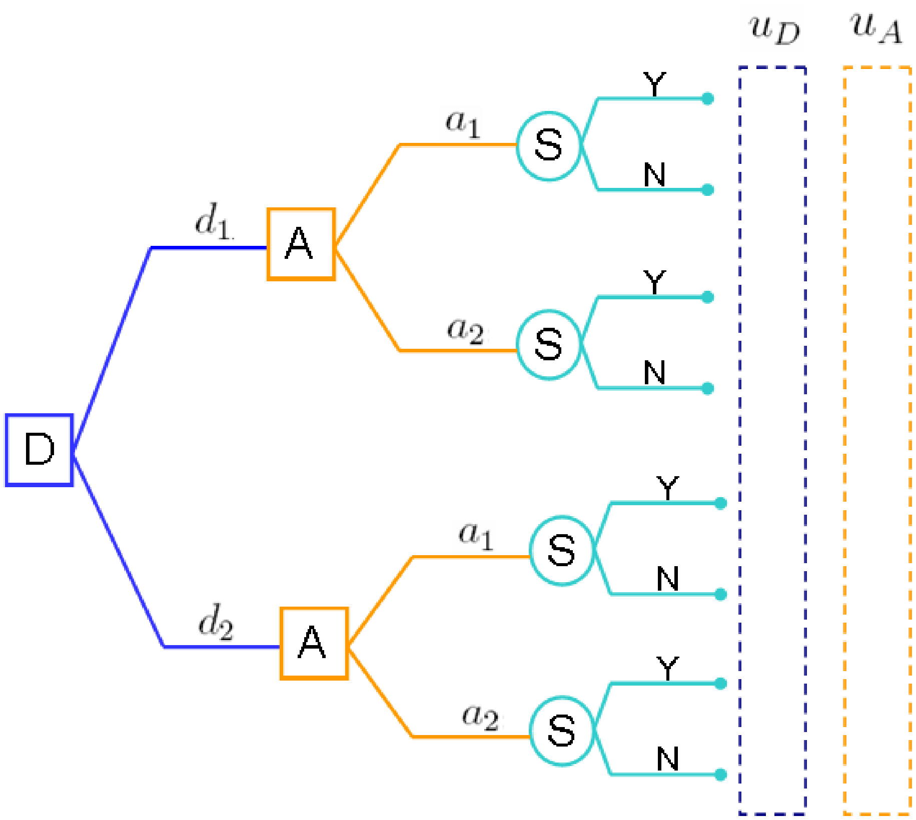

Standard Game Theory Analysis

Expected utilities at node S

$$\psi_D(d,a) = p_D(S=0|d,a)\; u_D(d,S=0) \;+\; p_D(S=1|d,a)\; u_D(d,S=1)$$

$$\psi_A(d,a) = p_A(S=0 \mid d,a)\; u_A(a,S=0) \;+\; p_A(S=1 \mid d,a)\; u_A(a,S=1)$$

Best Attacker's decision at node A

$$a^*(d) = \mathrm{argmax}_{a \in \mathcal{A}} \psi_A(d,a)$$

Assuming Defender knows Attacker's analysis,
Defender's best decision at node D

$$d^* = \mathrm{argmax}_{d \in \mathcal{D}} \psi_D(d, a^*(d))$$

Nash Solution: $(d^*, a^*(d^*))$

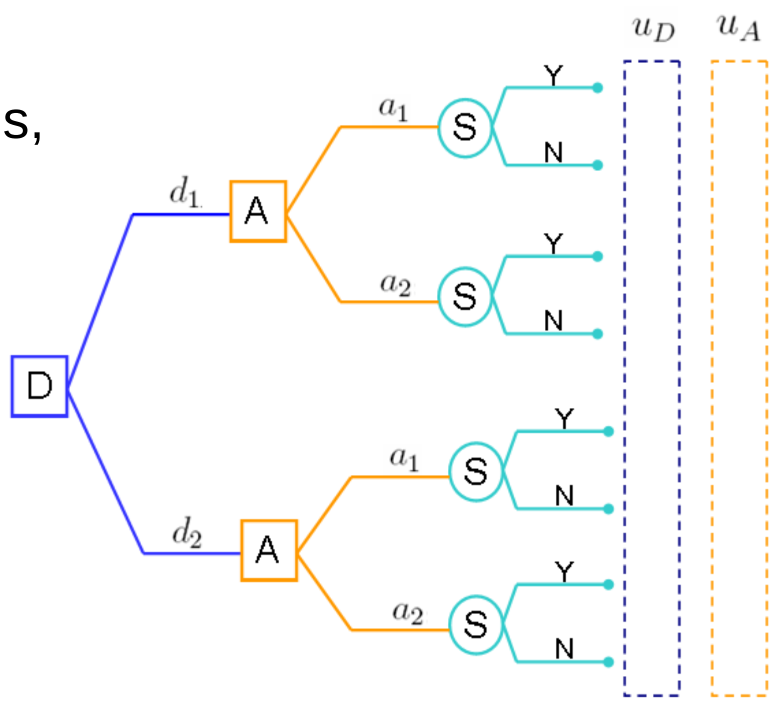

Supporting the Defender

Defender's problem	Defender's solution of maximum SEU
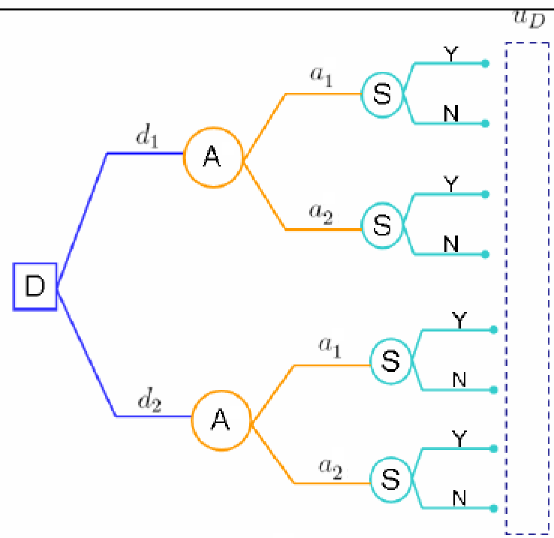 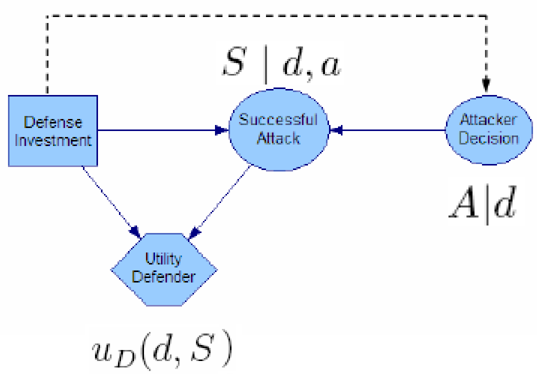	$\psi_D(d,a) = p_D(S=0\vert d,a)\ u_D(d,S=0)\ + p_D(S=1\vert d,a)\ u_D(d,S=1)$ $\psi_D(d) = \psi_D(d,a_1)\ p_D(A=a_1\vert d) + \psi_D(d,a_2)\ p_D(A=a_2\vert d)$ $d^* = \arg\max_{d \in X_D} \psi_D(d)$ Modeling input: $p_D(S\vert a,d)$ $p_D(A\vert d)$??

Supporting the Defender assessing Attacker's decision

Defender problem

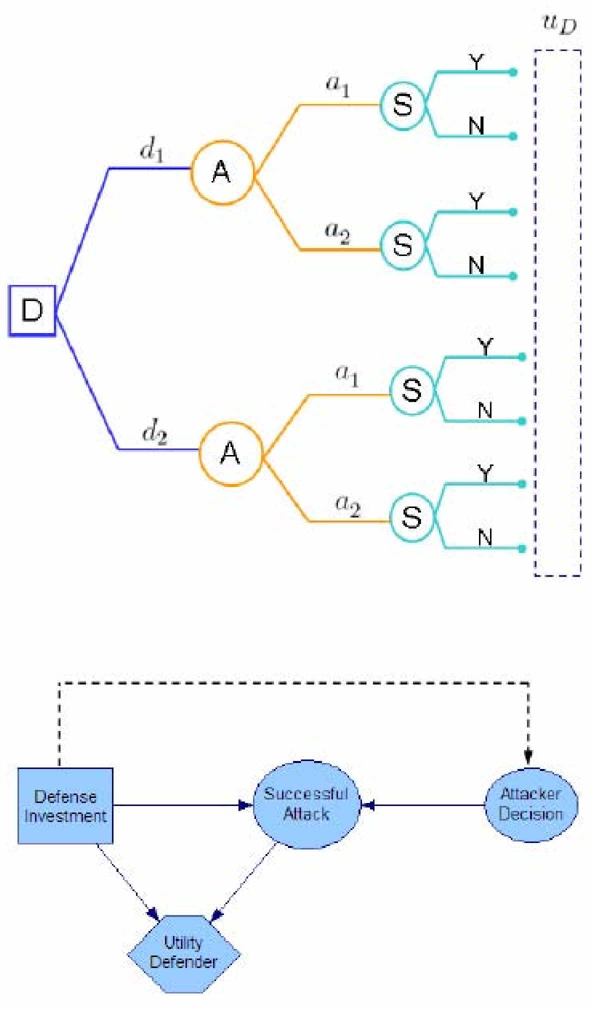

Defender's view of Attacker problem

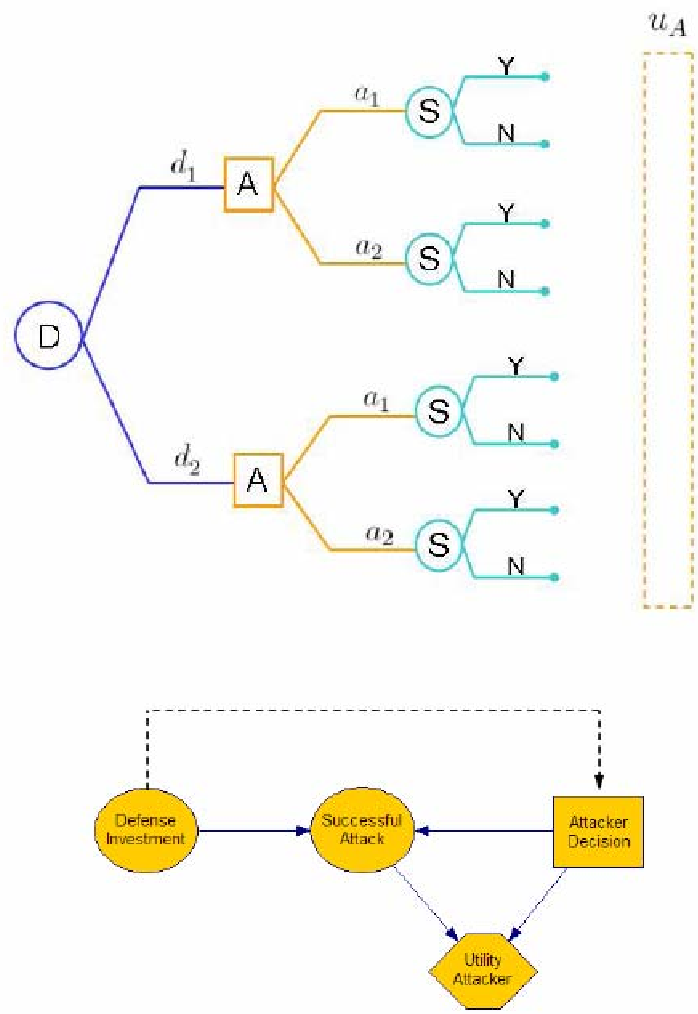

Solving the assessment problem

Defender's view of Attacker problem

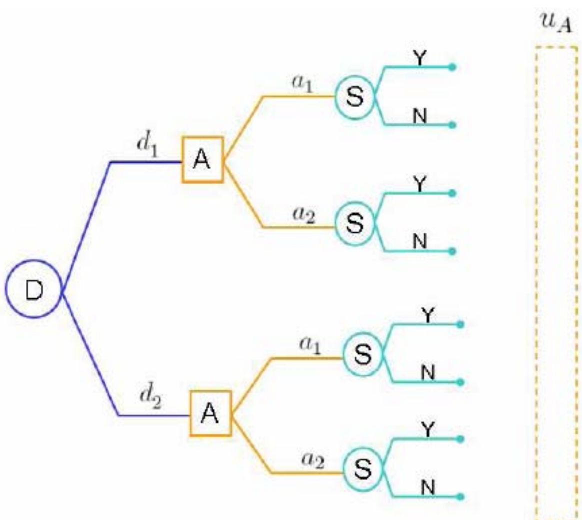

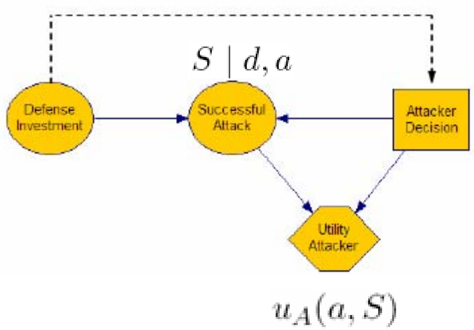

Elicitation of $p_D(A \mid d)$

A is a EU maximizer

D's beliefs about $(u_A, p_A) \sim (P_A, U_A) = F$

$$\Psi_A(d,a) = P_A(S=0 \mid d,a)\, U_A(a, S=0) \; + \\ P_A(S=1 \mid d,a)\, U_A(a, S=1)$$

$$p_D(A = a|d) = \mathbb{P}_F[a = \text{argmax}_{x \in \mathcal{A}} \Psi_A(d,x)]$$

<u>MC simulation</u>

$$\{(p_A^i, u_A^i)\}_{i=1}^n \sim F \longrightarrow \quad \psi_A^i \sim \Psi_A$$

$$a_i^*(d) = \text{argmax}_{x \in \mathcal{A}} \psi_A^i(x,d)$$

$$p_D(A = a|d) \approx \#\{a = a_i^*(d)\}/n$$

Bayesian solution for the Defend- Attack sequential model

1. Assess (p_D, u_D) from the Defender

2. Assess $F = (P_A, U_A)$, describing the Defender's uncertainty about (p_A, u_A)

3. For each d, simulate to assess $p_D(A|d)$ as follows:

   (a) Generate $(p_A^i, u_A^i) \sim F$, $i = 1, \ldots, n$

   (b) Solve $a_i^*(d) = \text{argmax}_{a \in \mathcal{A}} \psi_A^i(d, a)$

   (c) Approximate $\hat{p}_D(A = a|d) = \#\{a = a_i^*(d)\}/n$

4. Solve the Defender's problem

$$d^* = \text{argmax}_{d \in \mathcal{D}} \psi_D(d, a_1)\, \hat{p}_D(A = a_1|d) \;+\; \psi_D(d, a_2)\, \hat{p}_D(A = a_2|d)$$

How to avoid infinite regress?

$u_1 \quad p_1$	$u_2 \quad p_2$
$\hat{u}_2 \quad \hat{p}_2$	

G.T. (Full and common knowledge)

Asymmetric prescriptive/descriptive approach

$u_1 \quad p_1$	$u_2 \quad p_2$
$\hat{u}_2 \quad \hat{p}_2$	$\hat{u}_1 \quad \hat{p}_1$
$\hat{\hat{u}}_1 \quad \hat{\hat{p}}_1$	
$\vdots$	

Where to stop?

Discussion

- DA vs GT
  - A Bayesian prescriptive approach to support Defender against Attacker
  - Weaken common (prior) knowledge assumption
  - Analysis and assessment of Attacker’ thinking to anticipate their actions assuming Attacker is a expected utility maximizer
  - Computation of her defense of maximum expected utility
- The assessment problem under infinite regress
- Implementation issues
  - Elicitation of a valuable judgmental input from Defender
  - Computational issues

Outline

- Review of decision analysis
- Decision analytic framework for risk analysis.
- Aviation case study
- Adversarial risk analysis
- Some open problems

Some open problems

- DA, DAD, Simultaneous, Private information. What about general structures: Two Dis, with some shared nodes?
- Role of MCMC (Augmented probability simulation)
- Many unformalised criteria, very different in various fields. Could we unify them through decision theory, decision analysis?
- RA in ICT on decision theoretic footing
- Modeling ICT threats based on data
- Cyberinsurance
- 1 vs 1. m vs n
- What if the other is not EU maximiser?
- Case studies: e.g Pirates in Somalia
- ARA for other types of auctions