Adversarial Risk Analysis

Fabrizio Ruggeri

Istituto di Matematica Applicata e Tecnologie Informatiche Consiglio Nazionale delle Ricerche *Via Alfonso Corti 12, I-20133, Milano, Italy, European Union*

fabrizio@mi.imati.cnr.it

www.mi.imati.cnr.it/fabrizio

- Do not expect theorems! Sorry!
- You would have got them if I had decided to give a course on *Bayesian robustness*, my major area of interest about theoretical aspects since late 80's, but I thought this topic was too narrow (and not so exciting, for you) for a 9 hours course
- An alternative topic could have been *Bayesian Analysis of Stochastic Processes* (mostly Markov chains/processes, Poisson processes, Queues, Reliability), on which I coauthored a book with D. Rios Insua and M. Wiper, and my longstanding topic for the Ph.D. courses I gave around the world, but I saw that parts of it have been covered in a BIMSA/YMSC course
- Another possible course could have been on *Reliability*, my major area of interest about applications, but it would have been something on quite traditional methods and approaches
- I do not know about the future, but it would be nice (surely for me!) if I could have the opportunity to talk about them another time and discuss possible projects in those areas

- Then, why Adversarial Risk Analysis?
- More than ten years ago I was involved in a year-long program at SAMSI (Durham. NC, USA) where D. Banks, J. Rios and D. Rios Insua started developing new ideas on Adversarial Risk Analysis (ARA) from a Bayesian perspective
- My involvement was very minor at that time but I started working more seriously on ARA during another SAMSI program I led in 2019-20
- Since then, I have been working on many different aspects of ARA with D. Rios Insua (Madrid, Spain) and R. Soyer (Washington, DC, USA)
- I never had time to do a very thorough study of ARA (tough life for a researcher!) but this course has given me (Thanks!) the opportunity to deepen my knowledge, especially through the book on which many slides are based upon:

David L. Banks, Jesus M. Rios Aliaga, David Rios Insua (2016). Adversarial Risk Analysis. CRC Press

Adversarial Risk Analysis



David L. Banks Jesus Rios David Ríos Insua Winner of the 2017 De Groot Prize

A CHAPMAN & HALL BOOK



- I started as a mathematician but then I specialised in Bayesian (Mathematical) Statistics, with interest also in Bayesian Decision Analysis
- The topic will be presented from a decision theoretic and inferential Bayesian point of view, and not from a game theoretic one (there will be a course shortly on Game Theory)
- What should you expect since there will be no theorems? Models, models, models!
 - Models about beliefs
 - Models about preferences
 - Models about conceptual reasoning
- All of this will be done using tools from Bayesian Inference and Decision Analysis (prior and posterior distributions, elicitation, utility/loss functions, subjective expected utility maximisation principle, etc.)
- It is a hot topic, probably unknown to most of you, which allows me to present ideas typical of the Bayesian approach

OUTLINE OF THE COURSE

- Introduction to Bayesian Statistics
- Introduction to ARA, Game Theory, prior elicitation, robustness, influence diagram
- Discrete simultaneous games and modelling opponents
- Sequential games
- Examples (Somali Pirates)
- My work on Adversarial Hypothesis Testing
- My work on Acceptance Sampling
- My work on Adversarial Classification
- My work on Adversarial Software Testing

ALL BAYESIANS IN DAILY LIFE?

Visit Valparaiso or not?

- Prior knowledge
 - Where is Valparaiso?
 - Exotic place in Chile
- Data collection
 - Hiking guide
 - Tour operator catalogue
 - City of Valparaiso official website

ALL BAYESIANS IN DAILY LIFE?

- Posterior knowledge
 - Probably not a good place for hiking
 - Probably no tour found in the catalogue
 - UNESCO World Heritage Site, Neruda, very important harbour (before Panama canal), funiculars, next to Viña del Mar (beaches and casino)
- Forecast:
 - Will I enjoy Valparaiso or not?
 - Cost and time to get there
- Decision: To go or not to go?
 - Interest in the place
 - Distance and cost for travel, lodging and meals
 - Spanish language

BAYES THEOREM

- Patient subject to medical diagnostic test (P or N) for a disease D
- Sensitivity .95, i.e. $\mathbb{P}(P|D) = .95$
- Specificity .9, i.e. $\mathbb{P}(P^C|D^C) = .9$
- Physician's belief on patient having the disease 1%, i.e. $\mathbb{P}(D) = .01$
- Positive test $\Rightarrow \mathbb{P}(D|P)$?

BAYES THEOREM

$$\mathbb{P}(D|P) = \frac{\mathbb{P}(D \cap P)}{\mathbb{P}(P)} = \frac{\mathbb{P}(P|D)\mathbb{P}(D)}{\mathbb{P}(P|D)\mathbb{P}(D) + \mathbb{P}(P|D^{C})\mathbb{P}(D^{C})}$$
$$= \frac{.95 \cdot .01}{.95 \cdot .01 + .1 \cdot .99} = .0875$$

Positive test updates belief on patient having the disease: from 1% to 8.75%

Prior opinion updated into posterior one

BAYES THEOREM

• Partition $\{A_1, \ldots, A_n\}$ of Ω and $B \subset \Omega : \mathbb{P}(B) > 0$

$$\mathbb{P}(A_i|B) = \frac{\mathbb{P}(B|A_i)P(A_i)}{\sum_{j=1}^n \mathbb{P}(B|A_j)P(A_j)}$$

• X r.v. with density $f(x|\lambda)$, prior $\pi(\lambda)$

$$\Rightarrow \text{posterior } \pi(\lambda|x) = \frac{f(x|\lambda)\pi(\lambda)}{\int f(x|\omega)\pi(\omega)d\omega}$$

BAYESIAN STATISTICS

Bayesian statistics is ...

- ... another way to make inference and forecast on population features (*practitioner's view*)
- ... a way to learn from experience and improve own knowledge (educated layman's view)
- ... a formal tool to combine prior knowledge and experiments (*mathematician's view*)
- ... cheating (*hardcore frequentist statistician's view*)

• . . .

NOTIONS OF PROBABILITY

- Classical (random choice, equally likely events)
- Frequentist (probability as asymptotic limit of frequency)
- Subjective/Bayesian
- Axiomatic (Kolmogorov), which contains the other three

Bayesian \Rightarrow need to specify subjective *P* in (Ω, \mathcal{F}, P)

T = person having a tumor in his/her life I = person having an infarction in his/her life

 $\mathbb{P}(T \cup I) = .2, \ \mathbb{P}(T) = .3, \ \mathbb{P}(I) = .05, \ \mathbb{P}(T \cap I) = .1$

T = person having a tumor in his/her life I = person having an infarction in his/her life

$$\mathbb{P}(T \cup I) = .2, \ \mathbb{P}(T) = .3, \ \mathbb{P}(I) = .05, \ \mathbb{P}(T \cap I) = .1$$

- $\mathbb{P}(T \cup I) \geq \mathbb{P}(T)$
- $\mathbb{P}(I) \geq \mathbb{P}(T \cap I)$

T = person having a tumor in his/her life I = person having an infarction in his/her life

 $\mathbb{P}(T \cup I) = .3, \ \mathbb{P}(T) = .2, \ \mathbb{P}(I) = .2, \ \mathbb{P}(T \cap I) = .15$

T= person having a tumor in his/her life I= person having an infarction in his/her life

 $\mathbb{P}(T \cup I) = .3, \ \mathbb{P}(T) = .2, \ \mathbb{P}(I) = .2, \ \mathbb{P}(T \cap I) = .15$

•
$$.3 = \mathbb{P}(T \cup I) = \mathbb{P}(T) + \mathbb{P}(I) - \mathbb{P}(T \cap I) = .25$$

• $\mathbb{P}(T \cup I) = .3$, $\mathbb{P}(T) = .2$, $\mathbb{P}(I) = .2$, $\mathbb{P}(T \cap I) = .1$

 \Rightarrow assessments should comply with probability rules

- P(A): Probability one of us was born on a given day, say May, 1st
- $n \text{ people} \Rightarrow P(A) = 1 (364/365)^n$

ullet

$$n = 10 \implies P(A) = 0.027$$

$$n = 50 \implies P(A) = 0.128$$

$$n = 100 \implies P(A) = 0.240$$

$$n = 200 \implies P(A) = 0.422$$

$$n = 300 \implies P(A) = 0.561$$

• Therefore, what is your opinion about P(A)?

ILLUSTRATIVE EXAMPLE: FREQUENTIST APPROACH

Light bulb lifetime $\Rightarrow X \sim \mathcal{E}(\lambda) \& f(x; \lambda) = \lambda e^{-\lambda x} \quad x, \lambda > 0$

- Sample $\underline{X} = (X_1, \dots, X_n)$, i.i.d. $\mathcal{E}(\lambda)$
- Likelihood $l_x(\lambda) = \prod_{i=1}^n f(X_i; \lambda) = \lambda^n e^{-\lambda \sum_{i=1}^n X_i}$
- MLE: $\hat{\lambda} = n / \sum_{i=1}^{n} X_i$, C.I., UMVUE, consistency, etc.

What about available prior information on light bulbs behaviour? How can we translate it? \Rightarrow model and **parameter**

ILLUSTRATIVE EXAMPLE: BAYESIAN APPROACH

Light bulb lifetime $\Rightarrow X \sim \mathcal{E}(\lambda) \& f(x; \lambda) = \lambda e^{-\lambda x} \quad x, \lambda > 0$

- Sample $\underline{X} = (X_1, \ldots, X_n)$, i.i.d. $\mathcal{E}(\lambda)$
- Likelihood $l_x(\lambda) = \prod_{i=1}^n f(X_i; \lambda) = \lambda^n e^{-\lambda \sum_{i=1}^n X_i}$
- Prior $\lambda \sim \mathcal{G}(\alpha, \beta), \pi(\lambda) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda}$
- Posterior $\pi(\lambda|\underline{X}) \propto \lambda^n e^{-\lambda \sum_{i=1}^n X_i} \cdot \lambda^{\alpha-1} e^{-\beta\lambda}$ $\Rightarrow \lambda|\underline{X} \sim \mathcal{G}(\alpha+n, \beta+\sum_{i=1}^n X_i)$

Posterior distribution fundamental in Bayesian analysis

PARAMETER ESTIMATION - DECISION ANALYSIS

- Loss function $L(\lambda, a)$, $a \in \mathcal{A}$ action space
- Minimise $\mathcal{E}^{\pi(\lambda|\underline{X})}L(\lambda, a) = \int L(\lambda, a)\pi(\lambda|\underline{X})d\lambda$ w.r.t. a
 - $\Rightarrow \hat{\lambda}$ Bayesian optimal estimator of λ
 - $\hat{\lambda}$ posterior median if $L(\lambda, a) = |\lambda a|$
 - $\hat{\lambda}$ posterior mean $\mathcal{E}^{\pi(\lambda|\underline{X})}\lambda$ if $L(\lambda, a) = (\lambda a)^2$

$$\mathcal{E}^{\pi(\lambda|\underline{X})}L(\lambda,a) = \int (\lambda-a)^2 \pi(\lambda|\underline{X}) d\lambda$$

= $\int \lambda^2 \pi(\lambda|\underline{X}) d\lambda - 2a \int \lambda \pi(\lambda|\underline{X}) d\lambda + a^2 \cdot 1$
= $\int \lambda^2 \pi(\lambda|\underline{X}) d\lambda - 2a \mathcal{E}^{\pi(\lambda|\underline{X})} \lambda + a^2$

PARAMETER ESTIMATION

• Light bulb: posterior mean $\hat{\lambda} = \frac{\alpha + n}{\beta + \sum_{i=1}^{n} X_i}$ \Rightarrow compare with

- prior mean
$$\frac{\alpha}{\beta}$$

- MLE
$$\frac{n}{\sum_{i=1}^{n} X_i}$$

• MAP (Maximum a posteriori or posterior mode) $\Rightarrow \hat{\lambda} = \frac{\alpha + n - 1}{\beta + \sum X_i}$

PRIOR AND DATA INFLUENCE

• Posterior mean:
$$\hat{\lambda} = \frac{\alpha + n}{\beta + \sum X_i}$$

• Prior mean: $\hat{\lambda}_P = \frac{\alpha}{\beta}$ (and variance $\sigma^2 = \frac{\alpha}{\beta^2}$)

• MLE:
$$\hat{\lambda}_M = n / \sum X_i$$

• $\alpha_1 = k\alpha$ and $\beta_1 = k\beta \Rightarrow \hat{\lambda}_{1P} = \hat{\lambda}_P$ and $\sigma_1^2 = \sigma^2/k$

• Posterior mean:
$$\hat{\lambda} = \frac{k\alpha + n}{k\beta + \sum X_i}$$

- $k \to 0 \Rightarrow$ prior variance $\to \infty \Rightarrow \hat{\lambda} \to n / \sum X_i$, i.e. MLE (prior does not count)
- $k \to \infty \Rightarrow$ prior variance $\to 0 \Rightarrow \hat{\lambda} \to \hat{\lambda}_P$, i.e. prior mean (data do not count)

•
$$n \to \infty \Rightarrow \hat{\lambda} \sim \frac{n}{\sum X_i}$$
, i.e. MLE (prior does not count)

PARAMETER ESTIMATION

Prior influence (multinomial data and Dirichlet prior)

$$(n_1, \ldots, n_k) \sim \mathcal{MN}(n, p_1, \ldots, p_k)$$

 $(p_1, \ldots, p_k) \sim \mathcal{D}ir(s\alpha_1, \ldots, s\alpha_k), \ \sum \alpha_i = 1, \ s > 0$

• Posterior mean:
$$p_i^* = \frac{s\alpha_i + n_i}{s+n}$$

• Prior mean:
$$\tilde{p}_i = \alpha_i$$

- MLE: $\frac{n_i}{n}$
- $\bullet \ s \to \mathbf{0} \Rightarrow p_i^* \to \mathsf{MLE}$
- $s \to \infty \Rightarrow p_i^* \to \tilde{p_i}$

CREDIBLE INTERVALS

- $\mathcal{P}(\lambda \in A | \underline{X})$, credible (and Highest Posterior Density) intervals
- Compare with confidence intervals
- Light bulb:

$$\mathcal{P}(\lambda \leq z | \underline{X}) = \int_0^z \frac{(\beta + \sum X_i)^{\alpha + n}}{\Gamma(\alpha + n)} \lambda^{\alpha + n - 1} e^{-(\beta + \sum X_i)\lambda} d\lambda$$

HYPOTHESIS TESTING

• One sided test: H_0 : $\lambda \leq \lambda_0$ vs. H_1 : $\lambda > \lambda_0$

 \Rightarrow Reject H_0 iff $\mathbb{P}(\lambda \leq \lambda_0 | \underline{X}) \leq \alpha, \alpha$ significance level

- Two sided test: H_0 : $\lambda = \lambda_0$ vs. H_1 : $\lambda \neq \lambda_0$
 - Do not reject if $\lambda_0 \in A$, $A \ 100(1 \alpha)\%$ credible interval
 - Consider $\mathbb{P}([\lambda_0 \epsilon, \lambda_0 + \epsilon] | \underline{X})$
 - Dirac measure: $\mathbb{P}(\lambda_0) > 0$ and consider $\mathbb{P}(\lambda_0 | \underline{X})$

PREDICTION

- Prediction $P(X_{n+1}|\underline{X}) = \int P(X_{n+1}|\lambda)\pi(\lambda|\underline{X})d\lambda$
- Light bulb: $X_{n+1}|\lambda \sim \mathcal{E}(\lambda), \ \lambda|\underline{X} \sim \mathcal{G}(\alpha + n, \beta + \sum X_i)$

•
$$f_{X_{n+1}}(x|\underline{X}) = (\alpha+n)\frac{(\beta+\sum X_i)^{\alpha+n}}{(\beta+\sum X_i+x)^{\alpha+n+1}}$$

MODEL SELECTION

Compare $\mathcal{M}_1 = \{f_1(x|\theta_1), \pi(\theta_1)\}$ and $\mathcal{M}_2 = \{f_2(x|\theta_2), \pi(\theta_2)\}$

• Bayes factor

$$\Rightarrow BF = \frac{f_1(x)}{f_2(x)} = \frac{\int f_1(x|\theta_1)\pi(\theta_1)d\theta_1}{\int f_2(x|\theta_2)\pi(\theta_2)d\theta_2}$$

BF	$2\log_{10}BF$	Evidence in favor of \mathcal{M}_1
1 to 3	0 to 2	Hardly worth commenting
3 to 20	2 to 6	Positive
20 to 150	6 to 10	Strong
> 150	> 10	Very strong

• Posterior odds

$$\Rightarrow \frac{P(\mathcal{M}_1|data)}{P(\mathcal{M}_2|data)} = \frac{P(data|\mathcal{M}_1)}{P(data|\mathcal{M}_2)} \cdot \frac{P(\mathcal{M}_1)}{P(\mathcal{M}_2)} = BF \cdot \frac{P(\mathcal{M}_1)}{P(\mathcal{M}_2)}$$

Choice of a prior

- Defined on suitable set (interval vs. positive real)
- Suitable functional form (monotone/unimodal, heavy/light tails, etc.)
- Mathematical convenience
- *Tradition* (e.g. lognormal for engineers)

Gamma prior - choice of hyperparameters

- $X_1, \ldots, X_n \sim \mathcal{E}(\lambda)$
- $f(X_1, \ldots, X_n | \lambda) = \lambda^n \exp\{-\lambda \sum X_i\}$
- $\lambda \sim \mathcal{G}(\alpha, \beta) \Rightarrow f(\lambda | \alpha, \beta) = \beta^{\alpha} \lambda^{\alpha 1} \exp\{-\beta \lambda\} / \Gamma(\alpha)$
- $\Rightarrow \lambda | X_1, \dots, X_n \sim \mathcal{G}(\alpha + n, \beta + \sum X_i)$

Gamma prior - choice of hyperparameters

•
$$\mathcal{E}\lambda = \mu = \alpha/\beta$$
 and $Var\lambda = \sigma^2 = \alpha/\beta^2$
 $\Rightarrow \alpha = \mu^2/\sigma^2$ and $\beta = \mu/\sigma^2$

- Two quantiles \Rightarrow (α , β) using, say, Wilson-Hilferty approximation. Third quantile specified to check consistency
- Hypothetical experiment: posterior $\mathcal{G}(\alpha + n, \beta + \sum X_i)$ $\Rightarrow \alpha$ sample size and β sample sum

Using data to choose hyperparameters

• choose a prior $\pi(\lambda|\omega)$ of given functional form and use data to fit ω , i.e. look for $\hat{\omega} = \underset{\omega \in \Omega}{\arg \max} f(data|\omega) = \underset{\omega \in \Omega}{\arg \max} \int f(data|\lambda)\pi(\lambda|\omega)d\lambda$ (empirical Bayes)

Typical example (hierarchical model)

- *i* batches of n_i light bulbs each
- light bulbs in same batch with same properties
- light bulbs in different batches with *similar* properties

Hierarchical model

- $X_{ij_i}|\lambda_i \sim \mathcal{E}(\lambda_i), i = 1, n, j_i = 1, n_i$
- $\lambda_i | \beta \sim \mathcal{E}(\beta)$, s.t. $\mathcal{E}\lambda_i = 1/\beta$
- "Pure" Bayesian approach \Rightarrow prior on β (more later)
- Empirical Bayes

$$\begin{aligned} -\lambda_i |\beta, \underline{d} \sim \mathcal{G}(n_i + 1, \beta + \sum x_{ij_i}), \lambda_i \perp \lambda_j | \underline{d} \\ f(\underline{d} | \beta) &= \int f(\underline{d} | \underline{\lambda}) \pi(\underline{\lambda} | \beta) d\underline{\lambda} \\ &= \int \prod_{i=1}^n \left\{ \lambda_i^{n_i} e^{-\lambda_i \sum x_{ij_i}} \cdot \beta e^{-\beta \lambda_i} \right\} d\underline{\lambda} \\ &= \beta^n \prod_{i=1}^n \left\{ \frac{n_i!}{(\beta + \sum x_{ij_i})^{n_i + 1}} \right\} \end{aligned}$$

- maximised by $\hat{\beta} \Rightarrow \lambda_i | \hat{\beta}, \underline{d} \sim \mathcal{G}(n_i + 1, \hat{\beta} + \sum x_{ij_i}), \forall i$

BAYESIAN SIMULATIONS

Alternative choice: $\lambda \sim \mathcal{LN}(\alpha, \beta)$

• no posterior in closed form \Rightarrow numerical simulation

Markov Chain Monte Carlo (MCMC):

- draw^(*) a sample $\lambda^{(1)}, \lambda^{(2)}, \dots$ (Monte Carlo) ...
- ... from a Markov Chain whose stationary distribution is ...
- ... the posterior $\pi(\lambda | \underline{X})$ and compute ...
- $\mathcal{E}(\lambda|\underline{X}) \approx \sum_{i=m+1}^{n} \lambda^{(i)}/(n-m)$, etc.

(*) For $\lambda = (\theta, \mu) \Rightarrow$ Gibbs sampler:

- draw $\theta^{(i)}$ from $\theta|\mu^{(i-1)}, \underline{X}|$
- draw $\mu^{(i)}$ from $\mu|\theta^{(i)}, \underline{X}$
- repeat until convergence

MCMC: REGRESSION

- $y = \beta_0 + \beta_1 x + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma^2)$
- $(y_1, x_1), \ldots, (y_n, x_n)$
- Likelihood $\propto (\sigma^2)^{-n/2} \exp\{\frac{1}{\sigma^2} \sum_{i=1}^n (y_i \beta_0 \beta_1 x_i)^2\}$
- Priors: $\beta_0 \sim \mathcal{N}, \beta_1 \sim \mathcal{N}, \sigma^2 \sim \mathcal{IG}$
- Full posterior conditionals:
 - $\beta_0 | \beta_1, \sigma^2 \sim \mathcal{N}$
 - $\beta_1 | \beta_0, \sigma^2 \sim \mathcal{N}$
 - $\sigma^2 | \beta_0, \beta_1 \sim \mathcal{IG}$
 - $\Rightarrow \mathsf{MCMC}$

ARA IN A NUTSHELL *

Adversarial risk analysis (ARA) is a relatively new area of research that informs decision-making when facing intelligent opponents and uncertain outcomes. It is a decision-theoretic alternative to classical game theory that uses Bayesian subjective distributions to model the goals, resources, beliefs, and reasoning of the opponent. It enables an analyst to express her Bayesian beliefs about an opponent's utilities, capabilities, probabilities and the type of strategic calculation that the opponent is using. Within that framework, the analyst then solves the problem from the perspective of the opponent while placing subjective probability distributions on all unknown quantities. This produces a distribution over the actions of the opponent that permits the analyst to maximise her expected utility, accounting for the uncertainty she has about the opponent.

*Based on Banks, Gallego, Naveiro, Rios Insua, 2020
ARA IN A NUTSHELL *

- Game theory is the standard approach to adversarial reasoning, and it has been applied, among many other areas, in politics, biology, economics, social sciences and cybersecurity. The cornerstone of game theory is the Nash equilibrium, in which no opponent can improve their outcome by any unilateral action.
- Nonetheless, the fundamental premises of game theory have been criticised and the main concerns are:
 - The classical formulation generally assumes that all participants in the game have the same beliefs about the other players, and that all players know those beliefs are known. This common knowledge assumption is frequently unrealistic. For example, in a three-person auction, it is quite possible for players A and B to have different distributions for the value to player C of the item on offer and that they will conceal that information.
 - The field of behavioural economics has repeatedly demonstrated that humans do not act as game theory would prescribe, so it is a poor predictor of real-world decisions.

*Based on Banks, Gallego, Naveiro, Rios Insua, 2020

ARA IN A NUTSHELL

Think of a football (soccer) game between D and A and you are the manager of the team D and your goal is to win the game, but you have to think also about the way the other manager is preparing the game and selecting the initial players and the strategy

- First of all, you have to think about the strategy of the opponent: he might decide to play a very defensive (offensive) game with a lot of defenders (attackers) in the initial squad or he might choose players at random (not caring about the high probability of being fired pretty soon ...) ⇒ Concept uncertainty
- You are not sure (although you have some guesses) about the preferences (utilities) of the opponent, i.e. if he prefers to play for a draw rather than playing very offensive to win the game but also with high chances of losing it. Furthermore, you do not know what he thinks about your decisions but, again, you could make some guess about it ⇒ Epistemic uncertainty
- Once you and the opponent have chosen the initial squads and the strategies, then there is uncertainty about the final result (in Italian we say "The ball is round", meaning that everything could happen) ⇒ Aleatory uncertainty

ARA IN A NUTSHELL *

One of the advantages of ARA is its ability to partition the uncertainty into three separate components:

- Aleatory uncertainty: uncertainty in the outcome, conditional on the choices of each the opponents, to be handled by conventional statistical risk analysis
- **Epistemic uncertainty**: uncertainty in the opponent's utility function and assessment of the probability of outcomes conditional on the decisions that are made (by the analyst and the opponent), to be handled by in a Bayesian framework, making subjective probability assessments about each of these quantities
- **Concept uncertainty**: uncertainty about how the opponent is making his decision, since he might be a game theorist and seeks an equilibrium solution, or, perhaps, he randomizes, or follows some other protocol

*Based on Banks, Gallego, Naveiro, Rios Insua, 2020

ARA IN A NUTSHELL *

To make this a little more concrete, consider a sealed bid auction between Daphne and Apollo, each of whom wants to own a first edition of the *Theory* of Games and Economic Behaviour. Daphne's aleatory uncertainty is the value she receives conditional on her bid and Apollo's. If she has not been allowed to examine the book prior to the auction, then its condition is a random variable-perhaps it is old and torn, or perhaps it has marginalia written by John Nash, and both circumstances affect its value. **Epistemic uncertainty** arises because Daphne does not know the value of the book to Apollo, nor what he thinks is the probability that he will win with a bid of x dollars, nor how much money Apollo has. The **concept uncertainty** reflects the fact that Daphne does not know whether Apollo is determining his bid using classical game theory, or whether he is simply bidding some unknown fraction of his true top-dollar value, or using some other principle.

*Cited from Banks, Gallego, Naveiro, Rios Insua, 2020

- Blotto Game: example of a two-person simultaneous finite deterministic zero-sum game
- Colonel Blotto has six battalions that he must allocate across three battlefields
- At least one battalion must be assigned to each location
- His opponent, Colonel Klink, controls six battalions and must also place at least one in each location
- Neither knows in advance how the other will assign his forces, but both know that, for each battlefield, the side with the larger number of battalions will win (and if both assign the same number to the same location, then there will be a draw)
- The winner of the Blotto game is the side that wins the majority of the battles

- Blotto Game: example of a two-person simultaneous finite deterministic zero-sum game
 - Two players moving simultaneously
 - Allocation of battalions is not known until the troops engage
 - The choice sets are finite: there is only a fixed number of ways that Colonel Blotto can allocate his troops, and similarly for Colonel Klink
 - The game is deterministic, since both opponents know how many battalions the other controls, and chance plays no role in the outcome at a battlefield (but that assumption could be relaxed)
 - The game is zero-sum because a win at a battlefield for Colonel Blotto is a loss for Colonel Klink, and vice versa
- The choice set for both colonels' allocations is the same, made of triplets summing up to 6 and no zeros:

$$(1,1,4)$$
 $(1,4,1)$ $(4,1,1)$ $(1,2,3)$ $(1,3,2)$
 $(2,1,3)$ $(2,3,1)$ $(3,1,2)$ $(3,2,1)$ $(2,2,2)$

- The allocations are just permutations of (1, 1, 4), (1, 2, 3) and (2, 2, 2)
- We consider only those permutations (Wikipedia presents a similar example of Blotto game as "the game in which two players each write down three positive integers in non-decreasing order and such that they add up to a pre-specified number *S*. Subsequently, the two players show each other their writings, and compare corresponding numbers")
- Of course, the move from all triplets to the three exemplifying permutations reduces the number of possible cases: in our reduced setup two pairs (1,2,3) give a draw but (1,2,3) against (2,3,1) leads to the defeat of the first player but he wins if the opponent chooses (3,1,2)
- Therefore, it is true that, on average, (1, 2, 3) in our setup leads to a draw

- Payoff matrix: Payoff is 1 for winning the majority of battlefields, -1 for losing the majority of battlefields, and 0 for draws
- The payoff matrix shows that (2, 2, 2) beats (1, 1, 4), and every other pair of choices yields a draw
- Colonel Blotto could choose (2,2,2) since no other choice can win, and he could win if Colonel Klink foolishly chooses (1,1,4)
- Colonel Blotto could o.w. choose (1,2,3) since he cannot lose if Colonel Klink plays (2,2,2) or (1,1,4), and, if Colonel Klink also plays (1,2,3), then a random assignment of his troop strength to specific battlefields implies that Colonel Blotto has 1/6 chance of winning ((1,2,3) vs. (3,1,2), 1/6 chance of losing ((1,2,3) vs. (2,3,1), and 2/3 chance of a draw (the other 4)

			Blotto	
		(1, 1, 4)	(1, 2, 3)	(2, 2, 2)
Klink	(1, 1, 4)	(0,0)	(0,0)	(1, -1)
	(1, 2, 3)	(0,0)	(0, 0)	(0, 0)
	(2, 2, 2)	$\left \begin{array}{c} (1,-1) \end{array} \right $	(0, 0)	(0, 0)

• Classical game theory sees both (2, 2, 2) and (1, 2, 3) as solutions

- Formally, a pair of choices is a **Nash equilibrium** if neither player can gain by unilaterally changing his choice
- This means that Colonel Blotto is making the best decision possible, taking account of Colonel Klink's decision, and symmetrically, Colonel Klink is making the best decision possible, taking account of Colonel Blotto's
- For the Blotto game, all four possible pairs of choices taken from $\{(2, 2, 2), (1, 2, 3)\}$ are Nash equilibria since, e.g. if Blotto chooses (2, 2, 2) and Klink (1, 2, 3) then the latter cannot move to (1, 1, 4) which would be favorable to Blotto (payoff 1) but unfavorable to himself (payoff -1)
- For two-person zero-sum games, von Neumann and Morgenstern (1944) proved that a Nash equilibrium solution always exists
- The game gets more complex as the number of battalions increases. When there are more than 12 battalions apiece, no pure strategy is a Nash equilibrium. For example, with 13 battalions, Colonel Blotto should choose allocation (3, 5, 5) with probability 1/3, allocation (3, 3, 7) with probability 1/3, and allocation (1, 5, 7) with probability 1/3, and Colonel Klink should do likewise

- The Blotto game is deliberately simplistic
- One such simplification concerns the payoff: a 1 for a win, a -1 for a loss, and a 0 for a draw
- In more realistic scenarios, the value of a win could be large (if it resolved the war) or small (if it were a minor skirmish)
- In game theory and decision analysis, one handles this valuation problem through the **utility** of an outcome, combining all the costs (human lives, financial resources, etc.) and benefits (final victory, promotion of the colonel, etc.) into two numbers that summarize the net payoff to Colonel Blotto and the net payoff to Colonel Klink
- A second simplification is the assumption that the outcome is deterministic, depending only upon the number of battalions that each opponent allocates
- By chance, an inferior force might defeat superior numbers, or force a draw

- Also, the cost of a defeat may be small, if an orderly retreat is achieved, or large, if there were a massacre
- Thus, it would be more realistic to describe the utility that is realised from a particular pair of allocations as a random variable, rather than some known quantity
- Realistic uncertainty causes other complications
 - It is unlikely that Colonel Blotto knows exactly the utility that Colonel Klink assigns to a win, loss, or draw
 - And Colonel Blotto may have received intelligence regarding the allocations Colonel Klink will make - he is not certain of the accuracy of the intelligence, but should it be ignored?
 - Finally, Colonel Blotto may not know if Colonel Klink is selecting his allocation based on the mathematical solution to a game theory problem, or whether he is using some other system
- In real life, all of these uncertainties are relevant to the problem and, typically, analysts attempt to express such uncertainty through probability distributions

- (Bayesian) Decision Analysis supports a Decision Maker (DM) in making decisions under uncertainty:
 - Set of alternatives (actions) $a \in \mathcal{A}$
 - Unknown parameter θ depending on state of nature
 - Consequence $c(a, \theta)$ of action a when θ occurs
 - Utility function $u(c(a, \theta))$
 - Posterior distribution $\pi(\theta|x)$ on parameter θ , after observing x
 - Optimal action satisfies the Maximum (Subjective) Expected Utility Principle:

$$a^* = \underset{a \in \mathcal{A}}{\operatorname{arg\,max}} \int u(c(a,\theta))\pi(\theta|x)d\theta$$

• Just one agent playing against Nature

- State of nature: $\theta = \{$ Rain today, No rain today $\}$
- Actions $a = \{$ stay at home, go out with umbrella, go out without umbrella $\}$
- Consequences $c(a, \theta)$, e.g., c(stay at home, No rain today) = fired at work or c(go out without umbrella, Rain today) = unable to meet an important customer
- Utility function $u(c(a, \theta))$, e.g., u(c(stay at home, No rain today)) = -100,000 (income loss, in euros, after being fired)
- Posterior distribution $\pi(\theta|x)$ on parameter θ , after observing x, e.g., rain in the previous days
- Optimal action (suppose *go out with umbrella*) satisfies the Maximum (Subjective) Expected Utility Principle:

$$a^* = \underset{a \in \mathcal{A}}{\operatorname{arg\,max}} \int u(c(a,\theta))\pi(\theta|x)d\theta$$

- We now frame the approach when there is an opponent, like in a two-person simultaneous game
- Suppose an Attacker (A) selects an action from the set $\mathcal{A} = \{a_1, \ldots, a_n\}$, e.g. bombing a train or an airport, and a Defender (D) chooses an action from the set $\mathcal{D} = \{d_1, \ldots, d_m\}$, e.g. more police at train station or airport
- For each pair of actions (a, d) there is a consequence $s \in S$, e.g. casualties or not
- We see the problem from the viewpoint of *D*
- $\pi_D(a)$: *D*'s belief about *A*'s probability of choosing action $a \in \mathcal{A}$
- $p_D(s|a,d)$: *D*'s subjective probability for each possible outcome $s \in S$ given every choice $(a,d) \in A \times D$
- $u_D(d, a, s)$: D's utility for each combination of outcome and pair of choices
- *D*'s expected utility maximised by choosing $d^* \in \mathcal{D}$ s.t.

$$d^* = \underset{d \in \mathcal{D}}{\operatorname{arg\,max}} \int_{s \in \mathcal{S}} \int_{a \in \mathcal{A}} u_D(d, a, s) p_D(s|a, d) \pi_D(a) dads$$

•
$$d^* = \underset{d \in \mathcal{D}}{\operatorname{arg\,max}} \int_{s \in \mathcal{S}} \int_{a \in \mathcal{A}} u_D(d, a, s) p_D(s|a, d) \pi_D(a) dads$$

- This is the mathematical formulation but how is it in practice?
- How do we choose $u_D(d, a, s), p_D(s|a, d)$ and $\pi_D(a)$?
- In a Bayesian framework the two probabilities could be either prior opinions based only on *D*'s expertise or come from a combination of past data and expertise (i.e. posterior probabilities)
- Now there will be a short excursus about two topics which are often neglected in Bayesian courses:
 - Prior elicitation
 - Consequences of uncertainty in prior/model/utility (aka Bayesian robustness or sensitivity)

From Uncertain Judgements, O'Hagan et al, Wiley

Fundamentals of Probability and Judgement

- Importance of the distinction between aleatory (i.e. randomness in the phenomenon) and epistemic (i.e. due to poor knowledge of the phenomenon) uncertainty. Elicitation focuses mostly on epistemic uncertainty but people are more comfortable in providing probability assessment about aleatory uncertainty
- Elicited probabilities are, in general, given as answers to questions by a facilitator and not pre-formed quantifications of pre-analyzed beliefs
- Sometimes experts are spreading probability uniformly over the entertained interval
- Elicited probabilities might be biased or incoherent. Nonetheless, the facilitator should represent expert's knowledge and beliefs as accurately as possible

The Elicitation Context

- Elicitation is best conducted as a face-to-face interaction between expert and facilitator
- Elicitation process has several, important stages
 - Background and preparation, i.e. identification of variables for which expert's assessment is needed
 - Identify and recruit experts, which understand the problem, have a good reputation, want to cooperate, are impartial and have no personal stake in the findings
 - Motivating and training the experts, explaining why the elicitation is conducted and presenting toy (elicitation) examples along with basic probability notions
 - Structuring and decomposing (i.e. understanding dependencies)
 - Elicitation (elicit summaries, fit distributions and assess adequacy)

The Psychology of Judgement under Uncertainty

- Research shows that humans cannot be guaranteed to act as rational agents who follow the prescriptions of probability and decision theory, choosing sometimes strate-gies to provide judgements which are not optimal
- Broad expertise in an area is not, *per se*, a guarantee of specification of coherent probabilities. The elicitation process has always to be preceded by training in probability and then assistance during it would be desirable
- The facilitator and the expert should be aware of the possibility of biases occurring during the elicitation and pay attention to avoid them
- The facilitator should recognize that biases could be inadvertently introduced during the elicitation process (e.g. due to a particular order of questions) and try to avoid them by properly structuring the process
- The level of details used to describe an event could affect its probability assessment. Furthermore, probabilities assessed for mutually exclusive events should be checked for coherency

The Psychology of Judgement under Uncertainty (continued)

- Several errors and biases can be attributed to adopting too narrow a focus: e.g. focussing too much on one instance of a broader set of events or on one hypothesis as opposed to considering alternative hypotheses
- It could be helpful to introduce procedures which induce the expert to think analytically, e.g. involving him/her in the coherence checks
- Past experiences and knowledge of past data by the expert should be taken in account
- Output of the elicitation process become input to the decision analytic model and information on possible biases and uncertainty could help the statistician to determine in which way perform a sensitivity analysis

The Elicitation of Probabilities

- Importance of feedback on training exercise to improve expert's ability in assessing probabilities
- Subjective probabilities should be well calibrated but they often are not: it should be considered in training and in pursuing sensitivity analysis when making decisions
- Interpretation of verbal expressions of uncertainty varies considerably across individuals and situations: attention should be paid!
- Alternative ways of describing could lead to different assessments, as a result of a different information-processing strategy
- Aids and procedures for debiasing experts' opinions should be implemented but experts may be reluctant to their use: in any case, it would be ethical to inform them before their use

Eliciting Distributions - General

- Eliciting a distribution implies getting a finite, *small* number of summaries from the expert, fit a distribution and then check its adequacy
- Elicitation of univariate distributions is usually done through summaries based on probabilities (e.g. individual probabilities, quantiles, credible intervals and, sometimes, ratios of probabilities like in medical statistics [e.g. odds ratios]). There is no common opinion about the best summaries, whereas HPD intervals are in general considered to be avoided
- There is some evidence that people prefer to assess probability ratios rather than probabilities but it is in general impractical to transform such statements into ones about probabilities
- Elicitation of multivariate distributions is more complex and it should start from identifying subsets of independent variables

Eliciting Distributions - General (continued)

- Approaches to assess association among variables include joint and conditional probabilities and regression relationship
- The expert can specify only a finite number of summaries and there exist many distributions compatible with such assessment; such aspect has to be reported and properly addressed (robustness)

Eliciting and Fitting a Parametric Distribution

- In complex situation a parametric model is chosen to represent expert knowledge and the elicitation process leads to estimation of the parameters
- Often modelling is done through conjugate priors
- Assumptions about the prior distributions should be checked with the expert and feedback and overfitting are recommended to avoid inaccurate priors
- feedback and overfitting can also highlight assessments that are out of line. They can be fixed by reassessment by the expert or through some form of averaging
- The most widely assessed summaries are central measures (mean. median or mode) and quantiles, whereas sometimes hypothetical samples are used. The mostly assessed quantiles are upper and lower quartiles (besides median) but the 0.33 and 0.67 quantiles are sometimes used since they lead to better calibrated distributions. Quantiles are usually elicited considering variable intervals

Eliciting and Fitting a Parametric Distribution (continued)

- Many elicitation methods in literature are just theoretical ones, never used in practice. It is better to use those developed for practical purposes and already used!
- Distributions more flexible than conjugate ones are available but, in general, no elicitation method has been proposed for them
- Interactive computing is important since it allows to determine future questions, to provide feedback to the expert and detect (and correct) possible incoherencies
- Development of a decent elicitation method requires some knowledge about psychology, about statistics (of course!) and computing (if an interactive method is involved)

Eliciting Distributions - Uncertainty and Imprecision

- The accuracy of any fitted distribution as representation of the expert's beliefs is compromised by the imprecision in the expert's stated summaries and by the fact that just a limited number of summaries can be elicited in practice
- Although there is no consensus on how to report such uncertainty (e.g. upper and lower bounds), it is important to acknowledge it when reporting a fitted distribution

Evaluating Elicitation

- Two forms of accuracy for an elicited prior distribution are that it should accurately reflect both the expert's opinion and reality. Whereas the latter can be easily measured, the former is paramount since it is the defining quality of a subjective distribution
- Poor performance in scoring rules (formula used to provide a measure of the accuracy of a set of judgements) is often considered, in psychological literature, as a symptom of poor elicitation but often it is due to inaccurate knowledge
- Scoring rules are useful for comparing assessors and assessment methods. They are also an incentive for expert to record their opinions well and are important for a good training
- Based upon frequency of use, probability scoring rule (also called quadratic or Briar score) is the preferred rule for assessing discrete distributions, whereas logarithmic scoring rules is the preferred one for continuous distributions

Evaluating Elicitation (continued)

- In principle various decompositions of the probability scoring rule could be used to focus on different features of a set of assessments, but their practical utility is unclear
- Very few empirical successes in relating coherence with better calibration (i.e. the process of comparing subjective probabilities for event outcomes with observed relative frequencies)
- Feedback is the best way to improve accuracy about a subjective distribution representing expert's opinions
- When possible, overfitting should be coupled with feedback
- Interactive software is almost essential for the effective use of feedback

Multiple Experts

- Simple average (equally-weighted linear opinion pool) of distributions from a number of experts provides a simple, robust, general method for aggregating expert knowledge.
- More complex mathematical aggregation have the potential to perform better but their success depends on the goodness of the elicitation process
- Group elicitation (face-to-face) has probably a greater potential since it should lead, through elicitation and discussions, to a shared distribution, exploiting all sources of knowledge. The result depends on the ability of the facilitator in
 - encouraging sharing of knowledge
 - encouraging recognition of expertise
 - studying and making use of feedback
 - avoiding the group being dominated by someone's over-strong opinions
 - avoiding individual and group (in general due overconfidence) biases

BAYESIAN ROBUSTNESS

From Berger (1985)

- $X \sim \mathcal{N}(\theta, 1)$
- Expert's opinion on prior P: median at 0, quartiles at ± 1 , symmetric and unimodal
- \Rightarrow Possible priors include Cauchy C(0, 1) and Gaussian $\mathcal{N}(0, 2.19)$
- Interest in posterior mean $\mu^{C}(x)$ or $\mu^{N}(x)$

\overline{x}	0	1	2	4.5	10
$\mu^{C}(x)$	0	0.52	1.27	4.09	9.80
$\mu^N(x)$	0	0.69	1.37	3.09	6.87

- Decision strongly dependent on the choice of the prior for large x
- Alternative: Posterior median w.r.t. posterior mean

BAYESIAN ROBUSTNESS

- Practical impossibility of specifying priors exactly matching experts' knowledge
- Prior elicitation subject to uncertainty and, possibly, some degree of arbitrariness introduced by the analyst, e.g. the functional form of the distribution
- Uncertainty in the choice of priors modelled through a class of distribution (the same might apply for loss functions and statistical models/likelihoods)
- Use of indices to measure the consequences (i.e. perform robustness analysis) of the choice of a class of priors on the quantities of interest (e.g. posterior mean)
- An answer to the criticism about the arbitrariness in the choice of the prior and a possible excessive influence

BAYESIAN ROBUSTNESS

Interest mostly on sensitivity to changes in the prior

- Choice of a class Γ of priors
- Computation of a robustness measure, e.g. range $\delta = \overline{\rho} \underline{\rho}$ $(\overline{\rho} = \sup_{P \in \Gamma} E_{P^*}[h(\theta)] \text{ and } \underline{\rho} = \inf_{P \in \Gamma} E_{P^*}[h(\theta)])$
 - δ "small" \Rightarrow robustness
 - δ "large", ${\sf \Gamma}_1 \subset {\sf \Gamma}$ and/or new data
 - δ "large", Γ and same data

CLASSES OF PRIORS

- $\Gamma_P = \{P : p(\theta; \omega), \omega \in \Omega\}$ (Parametric class) - $\Gamma_P = \{\mathcal{G}(\alpha, \beta) : \alpha/\beta = \mu\}$
- $\Gamma_Q = \{P : \alpha_i \leq P(I_i) \leq \beta_i, i = 1, \dots, m\}$ (Quantile class)
- $\Gamma_{GM} = \{P : \int h_i(\theta) dP(\theta) = a_i, i = 1, ..., m\}$ (Generalised moments class) - $h(\theta) = \int_{-\infty}^x f(t|\theta) dt \Rightarrow \int h(\theta) dP(\theta) = \int_{-\infty}^x f(t) dt$ (Prior predictive distribution)
- $\Gamma^{DB} = \{F \text{ c.d.} f. : F_l(\theta) \le F(\theta) \le F_u(\theta), \forall \theta\}$ (Distribution bounded class)
- $\Gamma_{\varepsilon} = \{P : P = (1 \varepsilon)P_0 + \varepsilon Q, Q \in \mathcal{Q}\}$ (ε -contamination class)
- $K_g = \{P : \varphi_P(x) \ge g(x), \forall x \in [0, 1]\}, g \text{ nondecreasing, continuous, convex:} g(0) = 0 \text{ and } g(1) \le 1 \text{ (Concentration function class)}$

CLASSES OF MODELS

 $0 \leq M(\cdot) \leq U(\cdot)$ given and *l* likelihood

- $\Gamma_{\epsilon} = \{f : f(x|\theta) = (1-\epsilon)f_0(x|\theta) + (1-\epsilon)g(x|\theta), g \in \mathcal{G}\}\$ (ϵ -contaminations)
- $\Gamma_{DR} = \{f : \exists \alpha \text{ s.t. } M(x \theta_0) \le \alpha f(x|\theta_0) \le U(x \theta_0) \forall x\}$ (density ratio class)
- $\Gamma_L = \{l : M(\theta) \le l(\theta) \le U(\theta)\}$ (likelihood neighbourhood)
- $f(x|\theta) \propto \omega(x) f_0(x|\theta), \omega \in \Omega$
- $\Omega = \{ \omega : \omega_1(x) \le \omega(x) \le \omega_2(x) \}$

CLASSES OF LOSSES

• Parametric classes $\mathcal{L}_{\omega} = \{L = L_{\omega}, \omega \in \Omega\}$

- $L(a,\theta) = \beta(\exp\{\alpha(a-\theta)\} - \alpha(a-\theta) - 1), \alpha \neq 0, \beta > 0$ (LINEX)

- $\mathcal{L}_U = \{L : L(\theta, a) = L(|\theta a|), L(\cdot) \text{ any nondecreasing function}\}$ (Hwang's universal class)
- $\mathcal{L}_{\epsilon} = \{L : L(\theta, a) = (1 \epsilon)L_0(\theta, a) + \epsilon M(\theta, a), M \in \mathcal{W}\}$ (ϵ -contamination class)
- $\Omega = \{L : L(\theta, a) = \int_{\Lambda} L_{\lambda}(\theta, a) dG(\lambda)\}$ (Mixtures of convex loss functions)
 - $L_{\lambda} \in \Psi$, family of convex loss functions, $\lambda \in \Lambda$
 - $G \in \mathcal{P}$, class of all probability measures on (Λ, \mathcal{A})

TOWARDS ARA: RISK ANALYSIS

- Risk analysis: A systematic analytical process for assessing, managing and communicating the risk performed to understand the nature of unwanted, negative consequences to human life, health, property or the environment (so as to reduce and eliminate it)
 - Risk assessment: Information on the extent and characteristics of the risk attributed to a hazard
 - Risk management: The activities undertaken to control the hazard
 - **Risk communication**: Exchange of info/opinions concerning risk and risk-related factors among risk assessors, risk managers and other interested parties
- Interest in
 - costs/losses/utilities
 - risky actions by nature or attacker
 - impact of actions and reactions by defender

TOWARDS ARA

- Risks might be produced by an intelligent adversary A
- Adversary A could be an expected utility maximiser
- Uncertainty about *A*'s probabilities and utilities
 - Model *A*'s decision problem
 - Assess *A*'s probabilities and utilities
 - Find A's action of maximum expected utility
- \Rightarrow Adversarial Risk Analysis
INFLUENCE DIAGRAMS

- An influence diagram is a graphical tool used to represent a decision problem
- It is a directed acyclic graph with three kinds of nodes:
 - decision nodes, shown as rectangles
 - chance (or uncertainty) nodes, shown as ovals
 - preference (or value) nodes, shown as hexagons
- The domains of the nodes are, respectively, all the possible decisions, values taken by random variables and utilities
- Arrows, or directed edges, between nodes describe the structure of the problem
 - An arrow that points to a chance node signifies that the distribution at that node is conditioned on the values of all nodes at its tail
 - An arrow that points to a preference node means that the utility function depends upon the values of all nodes at its tail
 - An arrow that points to a decision node means that the choice made at that node is selected with knowledge of the values of all nodes at its tail

INFLUENCE DIAGRAMS

(From an early homework)

- You are currently a student and you have the possibility of applying for a 5 days only job. If you get the job, then you will be paid 500\$. There is a 100\$ non refundable application fee for the job: if you do not want to apply, you do not pay it and your total income is 0, of course.
- If you decide to apply, you will be interviewed by a manager. You know your current clothes are not *professional*, so that you will have to decide if to buy a new dress/suit for 100\$ or not. You know that your chances of getting a job are fifty-fifty if you show up professionally dressed, whereas they are just 1 to 3 if you wear your usual t-shirt.
- Using an influence diagram, can you tell which decision is the best one?



- The payoff is not in an hexagon ...
- The arrow from *Application fee* to *New Dress* is different since the latter decision depends on the former (new dress only if applying)

INFLUENCE DIAGRAMS

Expected utilities (I write outcome(probability))

- Do not apply: 0 (1) \Rightarrow 0
- Apply and buy: 300(.5) or -200(.5) $\Rightarrow 50$
- Apply and do not buy: 400 (.25) or -100 (.75) $\Rightarrow 25$
- Apply and buy is the optimal solution (but with different utilities ...)

- Discrete two-persons simultaneous game between D and A
- $\mathcal{D} = \{d_1, \ldots, d_m\}$ actions by D
- $\mathcal{A} = \{a_1, \ldots, a_n\}$ actions by A
- $\mathcal{X} = \{(X_{ij}^D, X_{ij}^A)\} \ m \times n \text{ bimatrix with payoffs to } D \text{ and } A \text{ for pair of actions } (i, j)$
- When there are r > 2 players, the bimatrix representation generalizes to an r-dimensional array
- In most practical situations, the payoffs in the cells are not fixed numbers but rather random variables
- The two opponents often have different beliefs about the distributions of those random variables, and imperfect knowledge of what each other will do and achieve, e.g. *A* could attack successfully while *D* thinks the attack will probably fail
- Such situations violate the framework used in traditional game theory, especially the common knowledge assumption needed to implement the Nash equilibrium solution

- To a decision analyst, the bimatrix formulation is helpful because it distinguishes epistemic uncertainty (i.e., which row-column pair is chosen, given the selection of a specific solution concept) from aleatory uncertainty (the outcome from picking that row-column pair)
- Within any specific cell determined by the row–column choice, *D* can apply traditional probabilistic risk analysis methods based upon expert opinion, probability models, historical data, and so forth
- *D*'s analysis generates a distribution over the result when *D* and *A* choose that row–column pair of actions
- By combining that distribution with her own utility function, *D* can calculate the distribution for her payoff
- A similar analysis allows *D* to infer the distribution that *A* has for his payoff, and this enables deeper reasoning related to epistemic uncertainty

- For each pair of choices (d, a), D receives the utility $u_D(d, a, \omega)$ which depends upon both chosen actions and upon chance, as indicated by the random variable ω
- In problems with fixed, non-random payoffs, one omits ω
- *D*'s belief about the probability distribution for ω , conditional on the choice (d, a), is represented by $p_D(\omega|d, a)$
- Symmetrically, A receives the utility $u_A(d, a, \omega)$, and believes the conditional density of ω is $p_A(\omega|d, a)$
- D's expected utility, given the choices (d, a), is $\mathbf{E} [u_D(d, a, \omega)|d, a] = \int u_D(d, a, \omega) p_D(\omega|d, a) d\omega$
- Similarly, A's expected utility is $\int u_A(d,a,\omega) p_A(\omega|d,a) d\omega$

- From a practical viewpoint, it is simpler to first find the distributions of outcomes conditional on a specific pair of actions (*d*, *a*), and then find the corresponding utilities
- As an example, D could use the probability model $p_D(\omega|d, a)$ to describe her belief about the chance of not discovering a bomb, where d is D's allocation of policemen to trains and a is A's decision about which train to target
- Then, conditional on the outcome that the bomb is not discovered, *D* can separately assess her utility, which combines mortality, economic costs, and political capital
- $Y_D(d, a, \omega)$ and $u_D[Y_D(d, a, \omega)]$: D's random outcome and utility
- $Y_A(d, a, \omega)$ and $u_A[Y_A(d, a, \omega)]$: A's random outcome and utility
- $X_{ij}^D = u_D [Y_D(d_i, a_j, \omega)]$ and $X_{ij}^A = u_A [Y_A(d_i, a_j, \omega)]$



Multi-agent influence diagram (MAID) showing decision, chance, and utility nodes, together with shared information structure, for the simultaneous Defend-Attack problem

- To perform ARA one needs to address first concept uncertainty, then epistemic uncertainty and, finally, aleatory uncertainty
- Some common solution concepts are:
 - Non-strategic play, in which D believes that A will select an action without consideration of her choice, e.g. if A selects actions with probability proportional to the perceived utility of success or if A is a non-sentient opponent, such as a hurricane
 - Nash equilibrium, which implies that D believes A is assuming that he and D have a great deal of common knowledge
 - Level-k thinking, in which D believes A thinks k plies deep in an "I think that she thinks that I think ..." kind of reasoning. The level-0 case corresponds to non-strategic play
 - Mirroring equilibrium analysis, in which D believes A is modeling D's decision making in the same way that she is modeling his, and both use subjective distributions on all unknown quantities

- Usually, *D* does not know which solution concept *A* has chosen, but, based on previous experience with *A*, and perhaps input from informants or other sources, she can decide which solution concept *A* will choose or place a subjective probability distribution over his possible solution concepts
- In the former case *D* can move to model epistemic uncertainty
- In the latter case, *D* could then make the decision that maximises her expected utility against that weighted mixture of strategies
 - Each solution concept will lead (after handling the relevant epistemic and aleatory uncertainties) to a distribution over *A*'s actions
 - Then *D* weighs each distribution by her personal probability that *A* is using that solution concept
 - This generates a weighted distribution on A, A's action space, which reflects all of D's knowledge about the problem and all of her uncertainty
 - The approach is closely related to Bayesian model averaging

- The distinguishing feature of ARA is that it emphasizes the advantage of building a model for the strategic reasoning of an opponent
- Regarding epistemic uncertainty, this is handled differently for each solution concept that *D* thinks *A* might use
- For example, with the Nash equilibrium concept, *D* believes that *A* thinks they both know the same bimatrix of payoffs
- In that case, the relevant epistemic uncertainty is *D*'s distribution over the bimatrices that *A* may be using

- Regarding aleatory uncertainty, this concerns the non-strategic randomness in an outcome
- Given a particular row-column choice in the bimatrix, the payoffs to each party are usually stochastic
- In that case, *D* must assess her beliefs about the outcome probabilities, conditional on the row–column pair. It warrants emphasis that this is not the same as assessing her beliefs about what *A*'s distributions over those outcomes might be—that is a matter of epistemic uncertainty, since it requires her to model *A*'s reasoning
- Aleatory uncertainty can be addressed through traditional probabilistic risk analysis
- *D*'s beliefs should be informed by expert judgment, previous history, and appropriate elicitation methods
- Regrettably, risk analysis may be imprecise: experts are overconfident, previous history may be only partially relevant or even misleading, and the wide range of elicitation methods highlights the pitfalls in making complex judgments

Non-Strategic Analysis

- The simplest non-strategic game is one against a non-sentient opponent
- In that case, traditional risk analysis has long been accepted as the appropriate approach
- In such probabilistic risk analyses, the decision maker has a distribution over the kinds of events that may occur, and distributions over the costs of actions to mitigate or remedy the consequences
- All of these distributions reflect aleatory uncertainty
- Most decision analysts would agree that one should select the action that maximises expected utility
- Example: A ship sailing from Mumbai to Napoli could be threatened by Somali pirates. The ship captain might decide to go through Suez Canal (risking hijacking by pirates) or Cape of Good Hope (safe but longer).

- Off the coast of Somalia a ship is attacked with probability 0.005
- Conditional on an attack, the ship is successfully hijacked with probability 0.4
- Based on past attacks, it is known that the average ransom paid is €2.3M
- The additional costs for the ship owners is €0 if the ship is not hijacked when attempting the Suez Canal and €0.5M if going to Cape of Good Hope
- There could be different utility functions for money:
 - Risk neutral: $u_1(x) = x$
 - Constant absolute risk aversion (CARA): $u_2(x) = 1 \exp(-\alpha x)$
 - Hyperbolic absolute risk aversion (HARA): $u_3(x) = (x \alpha)^{1-\beta}/(1 \beta)$
- u₂ and u₃ are risk averse utilities, corresponding to people preferring a small guaranteed payoff to a random payoff that has larger expected value but some chance of being very small



- X random ransom with d.f. F uniform between \in 2M and \in 2.6M
- Expected utility for transit through Suez Canal: $0.005 \times 0.4 \times E_F [u(-X)] + 0.005 \times 0.6 \times u(0) + 0.995u(0)$
- Expected utility for transit through Cape of Good Hope: u(-0.5)

- Risk neutral: $u_1(x) = x$
- Constant absolute risk aversion (CARA): $u_2(x) = 1 \exp(-\alpha x)$
- Hyperbolic absolute risk aversion (HARA): $u_3(x) = (x \alpha)^{1-\beta}/(1 \beta)$

	Expected Utility	
Utility Function	Suez Canal	Cape of Good Hope
Risk Neutral, u_1	-0.005	-0.500
CARA, $\alpha = 0.5$	-0.004	-0.284
CARA, $\alpha = 2$	-0.211	-1.719
CARA, $\alpha = 4$	-24.929	-6.389
HARA, $\alpha = -3$, $\beta = 0.25$	3.035	2.651
HARA, $\alpha = -3$, $\beta = 0.5$	3.461	3.162

Route through Cape of Good Hope chosen only by a very risk averse captain

Nash equilibrium

- The minimax principle is the simplest example of the Nash equilibrium solution concept: minimise the maximum expected loss
- It is the mirroring of the maximin principle: maximise the minimum expected utility
- Example: A will either develop an anthrax attack or a smallpox attack, and D will stockpile either Cipro (against anthrax) or smallpox vaccine. Neither party has the capability to do both.
- We suppose that all the available budget has been allocated by both *D* and *A* so that the payoffs depend only on the number of deaths and survivors
- *D* models the payoff matrix for *A*, with her payoffs implicitly represented as the negative of *A*'s payoffs since we consider a zero-sum game

	Smallpox	Anthrax
Vaccine	W	Y
Cipro	X	Z

	Smallpox	Anthrax
Vaccine	-500	200
Cipro	100	-400

- If *D* knew *A*'s (*W*, *X*, *Y*, *Z*) values, she could apply the maximin principle to solve the game and discover the action *A* would choose, enabling her to make the best response
- We suppose that Anthrax can kill more people (200) than Smallpox (100) if no counteraction is taken
- In a population of 500 people we suppose that Vaccine protects all of them under a Smallpox attack but only 400 when Cipro is used under an Anthrax attack
- Under the minimax principle, A looks for the minimum payoff of his action (-400 if Anthrax and -500 if Smallpox) and then chooses the action (Anthrax) maximizing his minimum payoff
- At this point *D* chooses to invest in Cipro!

- Typically *D* will not know *A*'s payoff values
- Within ARA the payoffs can be considered by D as random variables (W, X, Y, Z) and a joint density f(w, x, y, z) would be specified, possibly based upon medical knowledge of the pathogens, military intelligence from informants, personal intuition, or all of these and more
- Eliciting joint probability distributions that combine information from multiple sources is non-trivial, but for now, assume that D has been able to specify f(w, x, y, z)
- Wrong solution of the maximin problem (but followed by some analysts): compute the expected values of W, X, Y and Z, and plug them in the payoff matrix (as before)
- The right way requires the computation of p^* , D's probability that A will attack with smallpox $(1 p^*)$ is the probability of an anthrax attack)

- $p^* = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{P}$ [smallpox attack|w, x, y, z] (w, x, y, z) dw dx dy dz
- To solve this integral, 24 disjoints regions of \mathbb{R}^4 , corresponding to different orderings of the r.v.'s (e.g. W < X < Y < Z), should be considered
- We suppose a continuous df so that the probability of ties is 0
- In each region a maximin problem should be solved to identify *A*'s action
- The maximin problem can be solved in the two-by-two example either by a fixed choice or a mixed strategy, with both attacks chosen according to some probabilities (more details in the book).
- Linear programming is needed for examples with larger tables/more players

	Smallpox	Anthrax
Vaccine	μ_{11}	μ_{12}
Cipro	μ_{21}	μ_{22}

- D can elicit her beliefs about the expected number of lives lost under each possible pair of choices (i, j), where i indicates her choice and j indicates A's
- *D*'s expected loss from stockpiling vaccine is $p^* \times \mu_{11} + (1 p^*) \times \mu_{12}$ and her expected loss from stockpiling Cipro is $p^* \times \mu_{21} + (1 p^*) \times \mu_{22}$
- *D* selects the action that minimises the expected number of deaths
- Note sometimes the interchange between *loss* and *utility*

Level-*k* thinking

- A level-k analysis allows one to model how deeply an opponent reasons about a game
- If *D* performs a level-1 analysis, she assumes that *A* is a level-0 thinker; i.e., his choice is non-strategic, and depends only upon his own payoffs or perhaps is made at random
- A level-2 analysis means that *D* believes that *A* is a level-1 thinker, who will model *D* as a level-0 thinker
- A level-3 analysis means that *D* believes that *A* is a level-2 thinker, and so forth
- In this framework, D wants to reason one level deeper than A

	Left	Right
Up	0, ?	10, ?
Down	10, ?	0, ?

- *D* as a level-0 thinker who does not know *A*'s payoffs and is not using ARA methods to place a subjective probability distribution over those payoffs
- The most common decision rules used in these situations are:
 - Minimax criterion, in which one minimises the largest possible loss (equivalent to the maximin rule, which maximises the smallest possible gain)
 - Minimax regret criterion, in which one minimises the maximum difference between the realised payoff and the best payoff that would have been possible
 - Hurwicz criterion, in which one maximises the weighted average of the best and worst payoffs associated to each alternative, with weight $\alpha \in [0, 1]$ given to the best payoff from each choice is called the optimism coefficient. This is equivalent to the minimax rule when $\alpha = 0$
 - Laplace criterion, in which one maximises the average payoff, considering all *A*'s choices as equiprobable

• For the previous matrix, none of those approaches can produce a clear recommendation and *D* must therefore choose arbitrarily

	Left	Right
Up	0,0	10, 10
Down	10, 0	0,10

- Now *D* is a level-1 thinker since she knows and uses *A*'s payoffs
- *A* is a level-0 thinker who chooses his best action (*Right*) without looking at *D*'s actions
- Given *A*'s choice, then *D* chooses *Up*
- Moving to higher levels implies more cumbersome computations

Mirroring equilibria

- Each player places a subjective distribution over the utilities and probabilities of the opponent
- *D*'s distributions should reflect her assumption that *A* is performing a similar analysis regarding her strategy. The term *mirroring* derives from this self-similar modeling of the opponent's decision making
- In practice, the analyst will be on D's side, helping her in modelling utilities and probabilities, as well as in guessing those by A
- Besides those assessments, it will be possible to model the D's guess about A's opinion about the optimal decision by D
- In this way, D can get a distribution on the optimal decisions by A and choose her optimal decision
- This will be more clear when looking at my work on ARA and next

- The Defender (D) chooses an action from the set $\mathcal{D} = \{d_1, \ldots, d_m\}$,
- The Attacker (A) selects an action from the set $\mathcal{A} = \{a_1, \ldots, a_n\}$
- For each pair of choices (d_i, a_j) , there is a common random variable ω which determines the utility $u_D(d_i, a_j, \omega)$ that D receives and the utility $u_A(d_i, a_j, \omega)$ that A receives
- Assume that D and A seek to maximise their expected utilities
- Given a pair of choices (d_i, a_j) , D believes that the density for ω is $p_D(\omega|d_i, a_j)$ and A believes it is $p_A(\omega|d_i, a_j)$
- Then D's and A's expected utilities for (d_i, a_j) are, respectively,
 - $\psi_D(d_i, a_j) = \int u_D(d_i, a_j, \omega) p_D(\omega | d_i, a_j) d\omega$
 - $\psi_A(d_i, a_j) = \int u_A(d_i, a_j, \omega) p_A(\omega | d_i, a_j) d\omega$

- It is possible to build a bimatrix $\{(\psi_D(d_i, a_j), \psi_A(d_i, a_j))\}$ of pairs of expected utilities
- If both players know the utility function and probability function of the other, and if they both know that these were common knowledge, then the values in the bimatrix can be used to compute Nash equilibria, typically leading to randomised strategies
- However, common knowledge does not hold in the applications considered here and so Nash equilibrium solutions are not applicable
- Without common knowledge, D will need to formulate a probability mass function $p_D(a)$ that represents her beliefs about the probabilities of A's choices
- Given that, D selects the action d^* that solves $\arg \max_{d \in D} \Psi_D(d)$, where

$$egin{aligned} \Psi_D(d) &=& \sum_{a\in\mathcal{A}}\psi_D(d_i,a)p_D(a) \ &=& \sum_{a\in\mathcal{A}}\left[\int u_D(d_i,a,\omega)p_D(\omega|d_i,a)d\omega
ight]p_D(a) \end{aligned}$$

100

- $d^* = \arg \max_{d \in D} \sum_{a \in \mathcal{A}} \left[\int u_D(d_i, a, \omega) p_D(\omega | d_i, a) d\omega \right] p_D(a)$
- D maximises her expected utility w.r.t. her distributions over ω and A's choice
- Suppose that *A* is non-strategic
 - If A is not choosing at random (e.g. A is Nature provoking hurricanes), then D will elicit $p_D(a)$ based on past data and/or expert opinion, e.g. on both occurrences and severity of hurricanes and costs and benefits of different hurricane protections
 - If A is choosing at random, then a Dirichlet-multinomial model can be used. If there are no historical data, then $p_D(a)$ could be a Dirichlet distribution with parameters $(\alpha_1, \ldots, \alpha_n)$, while historical data $(a_j \text{ chosen } x_j \text{ times}, j = 1, \ldots, n)$ are from a multinomial model, updating $p_D(a)$ into a $\mathcal{D}ir(\alpha_1 + x_1, \ldots, \alpha_n + x_n)$
 - A Dirichlet distribution $\mathcal{D}ir(\alpha_1, \ldots, \alpha_n)$ has density $\frac{\sum_{i=1}^n \Gamma(\alpha_i)}{\prod_{i=1}^n \Gamma(\alpha_i)} \prod_{i=1}^n x_i^{\alpha_i 1}$, with $\sum_{i=1}^n x_i = 1, x_i > 0$ and $\alpha_i > 0, i = 1, \ldots, n$

• When A is strategic, then D (usually) believes that he wants to maximise his expected utility, and seeks the action a^* that solves $\arg \max_{a \in A} \Psi_A(a)$, where

$$egin{array}{rll} \Psi_A(a) &=& \displaystyle\sum_{d\in\mathcal{D}}\psi_A(d,a_j)p_A(d) \ &=& \displaystyle\sum_{d\in\mathcal{D}}\left[\int u_A(d,a_j,\omega)p_A(\omega|d,a_j)d\omega
ight]p_A(d) \end{array}$$

- So A needs to find $p_A(d)$, his distribution over D's choice
- D does not know p_A(ω|d_i, a_j), u_A(d_i, a_j, ω) and p_A(d) but she can model her subjective beliefs about all three quantities through random probabilities and utilities {P_A(ω|d_i, a_j), U_A(d_i, a_j, ω), P_A(d)}
- D can solve her optimisation problem computing $p_D(a)$ through

$$A \sim rgmax_{a \in \mathcal{A}} \sum_{d \in \mathcal{D}} \left[\int U_A(d, a, \omega) P_A(\omega | d, a) d\omega \right] P_A(d)$$

•
$$A \sim \underset{a \in \mathcal{A}}{\operatorname{arg\,max}} \sum_{d \in \mathcal{D}} \left[\int U_A(d, a, \omega) P_A(\omega | d, a) d\omega \right] P_A(d)$$

- For k = 1, ..., K, a triplet $\left\{ p_A^{(k)}(\omega|d_i, a_j), u_A^{(k)}(d_i, a_j, \omega), p_A^{(k)}(d) \right\}$ is generated from the random triplet $\{ P_A(\omega|d_i, a_j), U_A(d_i, a_j, \omega), P_A(d) \}$ and the optimisation problem is solved for A, obtaining the optimal $A^{(k)}$
- Finally $p_D(a)$ is approximated by the empirical distribution of the $A^{(k)}$'s and D can solve her optimisation problem
- In developing the triplet $\{P_A(\omega|d_i, a_j), U_A(d_i, a_j, \omega), P_A(d)\}$, the first two components are usually easier to specify
- $P_A(\omega|d_i, a_j)$ does not involve strategy since it is just what *D* thinks is *A*'s belief about the distribution of the outcome when *D* selects d_i and *A* selects a_j
- Similarly, the uncertainty about *A*'s true utility function, $u_A(d_i, a_j, \omega)$, is often small since *D* has good information about *A*'s objectives and values, so $U_A(d_i, a_j, \omega)$ will have small dispersion

- The choice of $P_A(d)$, i.e. what A thinks about D's action, is difficult
- If A is looking for a Nash equilibrium solution, then D would need p_D and u_D , the probabilities and utilities that A ascribes to her
- *D* specifies not only the random (P_A, U_A) as before but also the random (P_D, U_D) , i.e. what *D* thinks about what *A* thinks about *D*
- For k = 1, ..., K, $\left\{ p_A^{(k)}(\omega | d_i, a_j), u_A^{(k)}(d_i, a_j, \omega), p_D^{(k)}(\omega | d_i, a_j), u_D^{(k)}(d_i, a_j, \omega) \right\}$ is generated from the random $\{ P_A(\omega | d_i, a_j), U_A(d_i, a_j, \omega), P_D(\omega | d_i, a_j), U_D(d_i, a_j, \omega) \}$
- For each pair (d_i, a_j) D computes the (random) expected utilities $(\Psi_D(d_i, a_j), \Psi_A(d_i, a_j))$

$$\Psi_D(d_i, a_j) = \int U_D(d_i, a_j, \omega) P_D(\omega | d_i, a_j) d\omega$$

$$\Psi_A(d_i, a_j) = \int U_A(d_i, a_j, \omega) P_A(\omega | d_i, a_j) d\omega$$

- For k = 1, ..., K, D obtains a bimatrix given by $\left(\psi_D^{(k)}(d_i, a_j), \psi_A^{(k)}(d_i, a_j)\right)$, for each pair (d_i, a_j) , and compute the Nash equilibria $(d^{(k)}, a^{(k)})$
- When there are multiple equilibria, *D* should give each equal weight to them
- As mentioned before, $p_D(a)$ is approximated by the empirical distribution given by $a^{(k)}$'s and D can solve her optimisation problem

$$d^* = \operatorname*{arg\,max}_{d \in \mathcal{D}} \sum_{a \in \mathcal{A}} \psi_D(d, a) p_D(a)$$

• Since in most cases there is no closed form solution, then *D* would have to use computational methods to estimate her best decision

- What we have just seen is an example of D acting as a level-2 thinker since she obtains $p_D(a)$, i.e. her opinion on A's action, considering A as a level-1 thinker who chooses his action supposing that D is a level-0 thinker, i.e. non-strategic
- Earlier we have seen the case of *D* acting as a level-1 thinker who assumes that *A* is non-strategic
- For D being a level-3 thinker, A should be thought as a level-2 thinker and then apply to him what we have presented for D as level-2 thinker
- And so forth ...