

# 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*

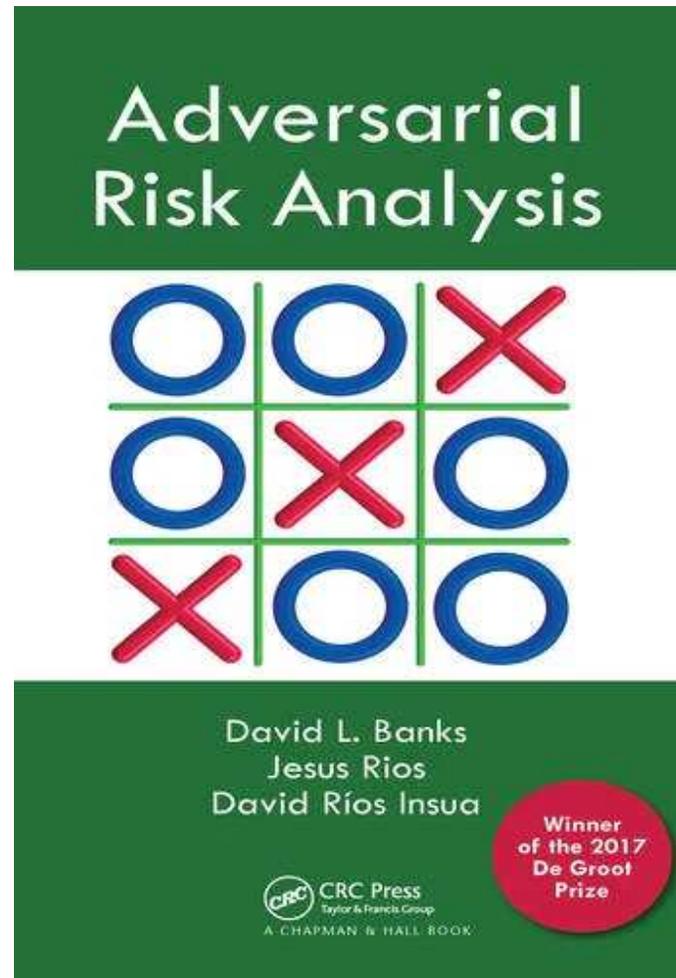
*fabrizio@mi.imati.cnr.it*

*www.mi.imati.cnr.it/fabrizio*

# OUTLINE OF THE COURSE

- Introduction to Bayesian Statistics
- Introduction to Adversarial Risk Analysis
- Discrete Simultaneous Games and Modelling Opponents
- Example: Auctions
- Sequential Games
- Example: Somali Pirates
- My works
  - Adversarial Hypothesis Testing
  - Batch Acceptance
  - Classification
  - Software Release

# REFERENCE BOOK



## 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 ...)  $\Rightarrow$  **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  $\Rightarrow$  **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)  $\Rightarrow$  **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

## QUICK GLIMPSE TO GAME THEORY

- 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

## QUICK GLIMPSE TO GAME THEORY

- 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)

## QUICK GLIMPSE TO GAME THEORY

- 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
  - Just consider all possible permutations for both colonels, assigning the same probability to each of the 36 pairs (6 permutations for each colonel)

## QUICK GLIMPSE TO GAME THEORY

- Payoff matrix: Payoff is 1 for winning the majority of battlefields,  $-1$  for losing the majority of battlefields, and 0 for draws
- $(x, y)$  in matrix:  $x$  payoff for Klink,  $y$  payoff for Blotto
- Matrix shows that  $(2, 2, 2)$  beats  $(1, 1, 4)$ , and every other pair 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	$(1, 1, 4)$	$(0, 0)$	$(0, 0)$	$(1, -1)$
	$(1, 2, 3)$	$(0, 0)$	$(0, 0)$	$(0, 0)$
	$(2, 2, 2)$	$(1, -1)$	$(0, 0)$	$(0, 0)$

## QUICK GLIMPSE TO GAME THEORY

- 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

## QUICK GLIMPSE TO GAME THEORY

- 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

## QUICK GLIMPSE TO GAME THEORY

- 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

## QUICK GLIMPSE TO GAME THEORY

- Game theory is the branch of mathematics that finds "optimal" solutions when two or more opponents with competing interests must select among possible actions
- There are important distinctions among games although some of them may not be so sharp. For example, the number of players can change over the course of a game, or games may evolve from competition to cooperation
- Opponents may make simultaneous or sequential decisions
  - Simultaneous: all players announce their decision at the same time
  - Sequential: players make their decisions over time, perhaps in response to previous decisions by other players. Most commonly, two players will alternate in declaring their decisions
- Games may be discrete or continuous
  - Discrete: countable (usually finite) set of choices, like in auctions where players must bid in euros
  - Continuous: uncountable set of choices, with possibility of bidding even arbitrary fractions of cents of euros

# QUICK GLIMPSE TO GAME THEORY

- Games may be deterministic or stochastic
  - Deterministic: all players know each other's possible actions and the consequences corresponding to every collection of choices
  - Stochastic: when the previous strong condition fails, e.g., when the payoff resulting from a specific combination of choices by the opponents is random
- Games may be zero-sum or non-zero-sum
  - Zero-sum: what one player gains another player must lose (e.g., as in gambling)
  - Non-zero-sum: the total gains or losses may be more or less than zero, e.g. when win-win solutions are possible among some or all the players (e.g., competing companies might collaborate to develop a product that creates new revenue for both)
- Games may or may not allow communication between the players
  - When communication is possible, then there is the potential for threats, bluffs, use of disinformation, and cooperation

# QUICK GLIMPSE TO GAME THEORY

- Games may have two or more players
  - When there are more than two players, then there is the possibility that some of them will form coalitions, which increases strategic complexity
- Games may be cooperative or non-cooperative
  - Cooperative: the players are able to form binding commitments externally enforced (e.g. through contract law)
  - Non-cooperative: the players cannot form alliances or all agreements need to be self-enforcing (e.g. through credible threats)
- Games may have perfect or imperfect information
  - Perfect: all players, at every move in the game, know the previous history of the game and the moves previously made by all other players (e.g. chess)
  - Imperfect: the players do not know all moves already made by the opponent such as a simultaneous move game (e.g. poker)

## QUICK GLIMPSE TO GAME THEORY

- Perfect information is often confused with complete information, which is a similar concept pertaining to the common knowledge of each player's sequence, strategies, and payoffs throughout game play
- Complete information requires that every player knows the strategies and payoffs available to the other players but not necessarily the actions taken, whereas perfect information is knowledge of all aspects of the game and players
- We just saw an example of a two-person simultaneous finite deterministic zero-sum game: the Blotto game

# QUICK GLIMPSE TO BAYESIAN DECISION ANALYSIS

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

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

- Just one agent playing against *Nature*

# QUICK GLIMPSE TO BAYESIAN DECISION ANALYSIS

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

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

## QUICK GLIMPSE TO BAYESIAN DECISION ANALYSIS

- 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  $\mathcal{A} = \{a_1, \dots, a_n\}$ , e.g. bombing a train or an airport, and a Defender ( $D$ ) chooses an action from the set  $\mathcal{D} = \{d_1, \dots, d_m\}$ , e.g. more police at train station or airport
- For each pair of actions  $(a, d)$  there is a consequence  $s \in \mathcal{S}$ , 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  $s \in \mathcal{S}$  given every choice  $(a, d) \in \mathcal{A} \times \mathcal{D}$
- $u_D(d, a, s)$ :  $D$ 's utility for each combination of outcome and pair of choices
- $D$ 's expected utility maximised by choosing  $d^* \in \mathcal{D}$  s.t.

$$d^* = \arg \max_{d \in \mathcal{D}} \int_{s \in \mathcal{S}} \int_{a \in \mathcal{A}} u_D(d, a, s) p_D(s|a, d) \pi_D(a) da ds$$

## QUICK GLIMPSE TO BAYESIAN DECISION ANALYSIS

- $d^* = \arg \max_{d \in \mathcal{D}} \int_{s \in \mathcal{S}} \int_{a \in \mathcal{A}} u_D(d, a, s) p_D(s|a, d) \pi_D(a) da ds$
- 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)
- In ARA 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

# 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

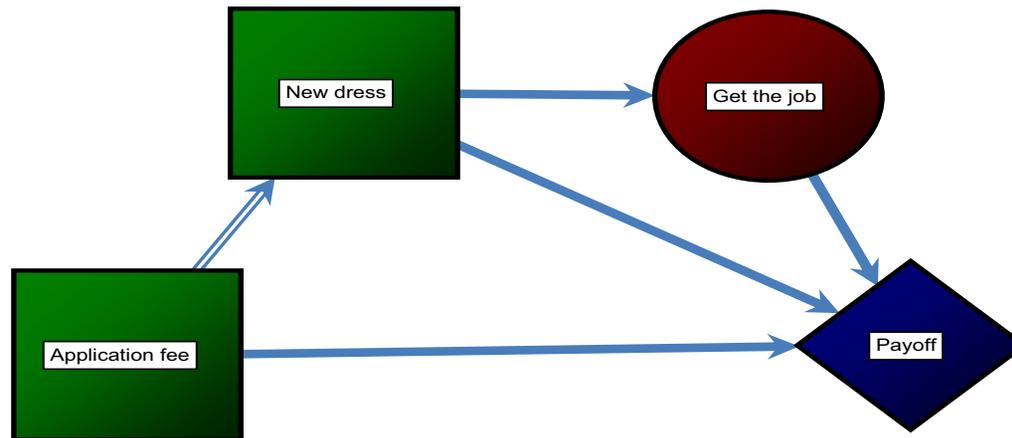
- An influence diagram is a graphical tool used to represent a decision problem
- It is a directed acyclic graph with three kinds of nodes:
  - decision nodes, shown as rectangles
  - chance (or uncertainty) nodes, shown as ovals
  - preference (or value) nodes, shown as hexagons
- The domains of the 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 means 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?

# INFLUENCE DIAGRAMS



- 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)*)

- $a_1$ : *Do not apply*:  $0\$(1) \Rightarrow 0$
- $a_2$ : *Apply and buy*:  $300\$(.5)$  or  $-200\$(.5) \Rightarrow 50$
- $a_3$ : *Apply and do not buy*:  $400\$(.25)$  or  $-100\$(.75) \Rightarrow 25$
- Decisions:  $d_0$  (job denied) and  $d_1$  (job offered)
- $U(a, d)$ : utility (here \$) when action  $a$  is taken and the decision is  $d$
- $P(d|a)$ : probability of a decision  $d$  when action  $a$  is taken
- Look for action maximising the expected utility:  
$$a^* = \arg \max_{a \in \{a_1, a_2, a_3\}} \{U(a, d_1)P(d_1|a) + U(a, d_0)P(d_0|a)\}$$
- *Apply and buy* is the optimal solution (but with different utilities ...)

# DISCRETE SIMULTANEOUS GAMES

- Discrete two-persons simultaneous game between  $D$  and  $A$
- $\mathcal{D} = \{d_1, \dots, d_m\}$  actions by  $D$
- $\mathcal{A} = \{a_1, \dots, a_n\}$  actions by  $A$
- $\mathcal{X} = \{(X_{ij}^D, X_{ij}^A)\}$   $m \times n$  bimatrix with payoffs to  $D$  and  $A$  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

# DISCRETE SIMULTANEOUS GAMES

- 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

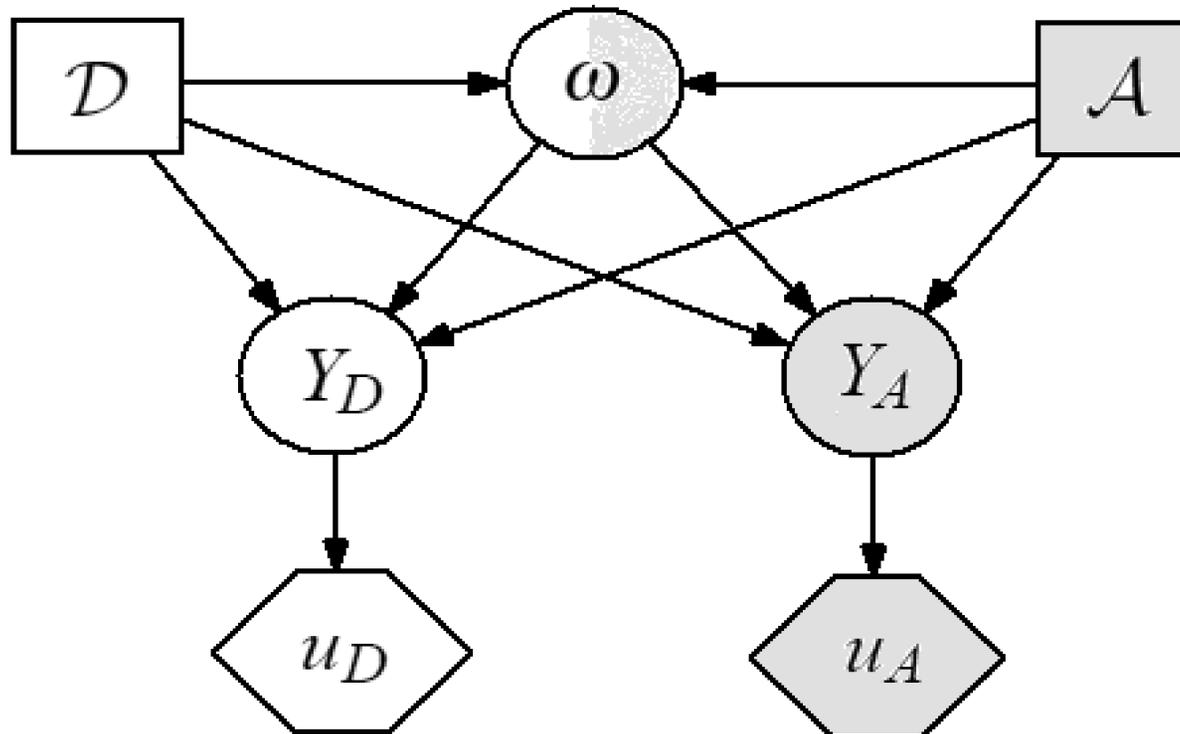
## DISCRETE SIMULTANEOUS GAMES

- 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  $\omega$
- In problems with fixed, non-random payoffs, one omits  $\omega$
- $D$ 's belief about the probability distribution for  $\omega$ , 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  $\omega$  is  $p_A(\omega|d, a)$
- $D$ 's expected utility, given the choices  $(d, a)$ , is
$$\mathbb{E} [u_D(d, a, \omega)|d, a] = \int u_D(d, a, \omega)p_D(\omega|d, a)d\omega$$
- Similarly,  $A$ 's expected utility is  $\int u_A(d, a, \omega)p_A(\omega|d, a)d\omega$

## DISCRETE SIMULTANEOUS GAMES

- 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  $X_{ij}^A = u_A [Y_A(d_i, a_j, \omega)]$

## DISCRETE SIMULTANEOUS GAMES



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

# DISCRETE SIMULTANEOUS GAMES

- 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

# DISCRETE SIMULTANEOUS GAMES

- 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  $\mathcal{A}$ ,  $A$ '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

## DISCRETE SIMULTANEOUS GAMES

- 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

## DISCRETE SIMULTANEOUS GAMES

- 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

# MODELLING OPPONENTS

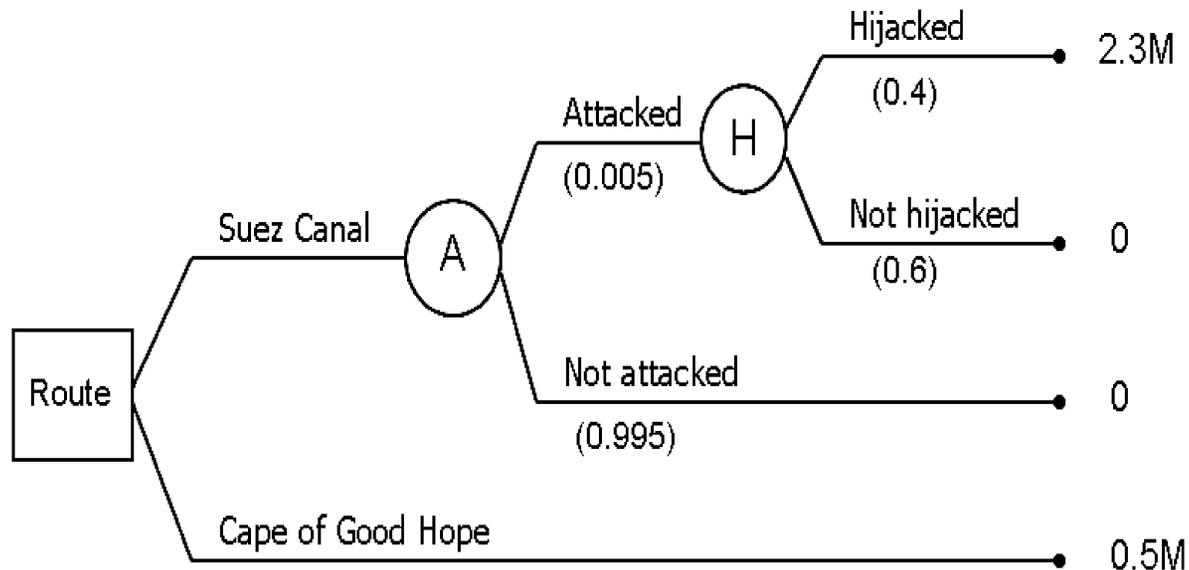
## Non-Strategic Analysis

- 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).

# MODELLING OPPONENTS

- 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):  $u_3(x) = (x - \alpha)^{1-\beta} / (1 - \beta)$
- $u_2$  and  $u_3$  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

# MODELLING OPPONENTS



- $X$  random ransom with d.f.  $F$  uniform between €2M and €2.6M
- Expected utility for transit through Suez Canal:  

$$0.005 \times 0.4 \times \mathbf{E}_F [u(-X)] + 0.005 \times 0.6 \times u(0) + 0.995u(0)$$
- Expected utility for transit through Cape of Good Hope:  $u(-0.5)$

## MODELLING OPPONENTS

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

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

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

# MODELLING OPPONENTS

## 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	$X$	$Z$

## MODELLING OPPONENTS

	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!

## MODELLING OPPONENTS

- 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)

## MODELLING OPPONENTS

- $p^* = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{P} [\text{smallpox attack} | w, x, y, z] f(w, x, y, z) dw dx dy dz$
- To solve this integral, 24 disjoint regions of  $\mathbf{R}^4$ , 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

## MODELLING OPPONENTS

	Smallpox	Anthrax
Vaccine	$\mu_{11}$	$\mu_{12}$
Cipro	$\mu_{21}$	$\mu_{22}$

- $D$  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*

# MODELLING OPPONENTS

## 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  $A$

## MODELLING OPPONENTS

	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

## MODELLING OPPONENTS

- 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

# MODELLING OPPONENTS

## 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

## ARA IN PRACTICE

- The Defender ( $D$ ) chooses an action from the set  $\mathcal{D} = \{d_1, \dots, d_m\}$ ,
- The Attacker ( $A$ ) selects an action from the set  $\mathcal{A} = \{a_1, \dots, a_n\}$
- For each pair of choices  $(d_i, a_j)$ , there is a common random variable  $\omega$  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  $\omega$  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,
  - $\psi_D(d_i, a_j) = \int u_D(d_i, a_j, \omega)p_D(\omega|d_i, a_j)d\omega$
  - $\psi_A(d_i, a_j) = \int u_A(d_i, a_j, \omega)p_A(\omega|d_i, a_j)d\omega$

## ARA IN PRACTICE

- 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  $\arg \max_{d \in \mathcal{D}} \Psi_D(d)$ , where

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

## ARA IN PRACTICE

- $d^* = \arg \max_{d \in \mathcal{D}} \sum_{a \in \mathcal{A}} \left[ \int u_D(d_i, a, \omega) p_D(\omega | d_i, a) d\omega \right] p_D(a)$
- $D$  maximises her expected utility w.r.t. her distributions over  $\omega$  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  $(\alpha_1, \dots, \alpha_n)$ , while historical data ( $a_j$  chosen  $x_j$  times,  $j = 1, \dots, n$ ) are from a multinomial model, updating  $p_D(a)$  into a  $\text{Dir}(\alpha_1 + x_1, \dots, \alpha_n + x_n)$
  - A Dirichlet distribution  $\text{Dir}(\alpha_1, \dots, \alpha_n)$  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  $\alpha_i > 0, i = 1, \dots, n$

## ARA IN PRACTICE

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

$$\begin{aligned} \Psi_A(a) &= \sum_{d \in \mathcal{D}} \psi_A(d, a_j) p_A(d) \\ &= \sum_{d \in \mathcal{D}} \left[ \int u_A(d, a_j, \omega) p_A(\omega | d, a_j) d\omega \right] p_A(d) \end{aligned}$$

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

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

## ARA IN PRACTICE

- $A \sim \arg \max_{a \in \mathcal{A}} \sum_{d \in \mathcal{D}} \left[ \int U_A(d, a, \omega) P_A(\omega|d, a) d\omega \right] P_A(d)$
- For  $k = 1, \dots, K$ , 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

## ARA IN PRACTICE

- 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  $k = 1, \dots, K$ ,  $\left\{ p_A^{(k)}(\omega|d_i, a_j), u_A^{(k)}(d_i, a_j, \omega), p_D^{(k)}(\omega|d_i, a_j), u_D^{(k)}(d_i, a_j, \omega) \right\}$  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  $(\Psi_D(d_i, a_j), \Psi_A(d_i, a_j))$

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

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

## ARA IN PRACTICE

- For  $k = 1, \dots, K$ ,  $D$  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

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

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

## ARA IN PRACTICE

- 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 ...