AUCTIONS

- We now consider continuous asymmetric, first-price, independent-value, sealed-bid auctions with risk-neutral bidders
  - *Continuity* means that the bids may take any value in an interval, in contrast with the discrete games treated so far
  - Symmetric auctions assume that the values opponents have for the item on offer are randomly drawn from the same (known) distribution, whereas *asymmetric* auctions allow different opponents to draw from different (known) distributions
  - In *first-price* auctions, the highest bidder wins, and pays the amount of that bid
  - The *independent-value* condition implies that the private value that one bidder has for an object is not influenced by the private value that other bidders have for that object
  - The *sealed-bid* condition ensures that whatever initial information a bidder has about the value distributions of his opponents does not change as the auction proceeds (in contrast with, say, an English auction, where opponents place increasingly higher bids until all but one has dropped out)
  - *Risk neutrality* implies that each bidder attempts to maximise his expected monetary profit

AUCTIONS

- Auctions of this kind are common, and are especially popular when the commercial value of the item on offer is difficult to determine

- In many sectors (e.g. public one in Italy) companies make sealed bids on federal contracts, and the lowest qualified bidder prevails. This problem is formally equivalent to the one we are considering here with the highest bid

- We consider an example in which a lady named Bonnie (B) is bidding against a man named Clyde (C) for a first edition of *The Theory of Games and Economic Behaviour* (the book by von Neumann and Morgenstern, 1944, is the first one about Game Theory and I saw a copy of the first edition sold over Internet at $15,000$ USD)

- Also, this auction does not have a reservation price, i.e. a secret lower bound set by the owner, so that the book will not be sold if no bid exceeds the reservation price

AUCTIONS

- Auctions easily illustrate aleatory, epistemic, and concept uncertainties
  - Aleatory uncertainty arises when B does not have full knowledge of the condition of the book (e.g. damaged or annotated by John Nash)
  - Epistemic uncertainty about the private value of the opposing bidder C can be due to, e.g., his better knowledge of the condition of the book or a sentimental value
  - Concept uncertainty occurs when B does not know which kind of strategic analysis C will perform when calculating his bid

AUCTIONS

- Suppose B believes that C is non-strategic, i.e. his rule to select his bid does not depend upon his analysis of B's situation

- If B has a distribution F over C's bid and x_0 is her true value for the book, then, under the assumption that her utility function for money is linear, she will maximise her expected utility in a first-price auction by bidding

  $x^* = \arg\max_{x \in \mathcal{R}^+} (x_0 - x) F(x)$

- The right-hand side is B's expected utility since $(x_0 - x)$ is her utility (profit) and $F(x)$ the probability that a bid x wins the auction (Y is C's bid)

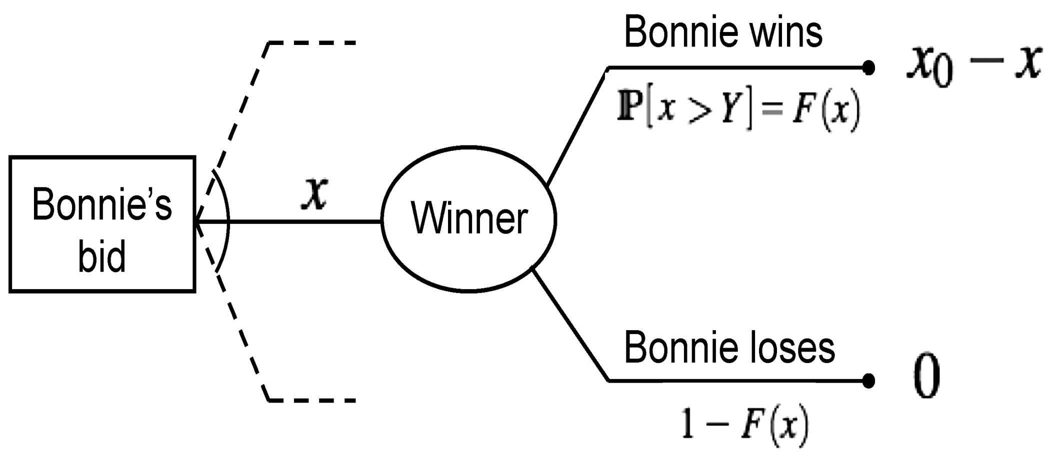

AUCTIONS

- Elicitation of $F(x)$ is based on B's knowledge, as discussed earlier
- An interesting approach is based on dividing the elicitation in two parts:
  - G_1: distribution over the value of the book to C
  - G_2: distribution on the fraction of his (i.e. for him!) true value that he bids
- G_1 could be based on past sales or appraisal values by experts, possibly adjusted if B thinks the book has a greater/lesser value for C
- G_2 could be based on knowledge of C's behaviour in past successful bids or past statements, or studies on underbidding

AUCTIONS

- V r.v. (random variable) with distribution G_1 (and density g_1) on $(0, \infty)$ denoting B's epistemic uncertainty about C's true value
- P r.v. with distribution G_2 (and density g_2) on $[0, 1]$ denoting the fraction of the true value bid by C
- $Y = PV$ r.v. with distribution F denoting the amount of C's bid
- We now assume that the true value and the proportional reduction are independent; if not, it is still possible to consider an integral, but more complex

$$
\begin{aligned}
F(y) = \mathbf{P}[PV \leq y] &= \int_0^y \int_0^1 g_1(v) g_2(p) dp dv + \int_y^\infty \int_0^{y/v} g_1(v) g_2(p) dp dv \\
&= G_1(y) + \int_y^\infty g_1(v) G_2(y/v) dv
\end{aligned}
$$

- The first term corresponds to a true value not exceeding y and the second to a true value exceeding y but still with a bid not larger than y

AUCTIONS

- Suppose B's personal value for the book on auction is $x_0 = \$150$

- B models C's value for the book as a uniform r.v. on $(0, 200)$ (in \$) with $G_1(v) = v/200, \ v \in [0, 200]$

- B models the fraction of C's value actually bid as a r.v. on $[0, 1]$ with $G_2(p) = p^9$

- $F(y) = G_1(y) + \int_y^\infty g_1(v) G_2(y/v) dv = \frac{9}{8}\frac{y}{200} - \frac{1}{8}\frac{y^9}{200^9}$, for $\$0 \le y \le \$200$

- $x^* = \arg\max_{x \in \mathcal{R}^+} (x_0 - x) F(x)$

  $\Rightarrow 0 = \frac{d}{dx}[(x_0 - x)F(x)] = 675 - 9x - \frac{675}{200^8}x^8 + \frac{5}{200^8}x^9$

- Numerical solution shows her bid should be about half of x_0, i.e. $x^* = \$75$

AUCTIONS

- Up to now there is neither concept uncertainty (since C is non-strategic, with bid supposed to be proportional to his true value) nor aleatory uncertainty (since B knows her value x_0)

- Suppose that x_0 is unknown and B's true value is modelled by a r.v. X_0 with distribution H and expected value μ

- $\Rightarrow$ Look for $x^* = \arg\max_{x \in \mathcal{R}^+} \mathbf{E}_H[(X_0 - x)F(x)] = (\mu - x)F(x)$

  (Remember that we are assuming $F(x)$ known!)

- Therefore, B should just elicit μ and not H: this is a consequence of the risk neutrality assumption, which implies that her utility for money is linear

- Note that sometimes aleatory uncertainty is more important than epistemic and concept ones, like for a company bidding for the construction of an oil plant, which does not know all the costs and difficulties that will arise, and thus does not know exactly what profit would be realised from its bid

AUCTIONS

- The maximin perspective here is not very helpful
- $u(x, y) = (x_0 - x)I_{(y,\infty)}(x)$
  - B's utility function when she bids x and C bids y
  - $I_A(\cdot)$ indicator function of the set A
  - We assume that B does not bid more than x_0
  - Utility equal to 0 if $x < y$ (i.e., C wins the bid) and positive if $y < x < x_0$ (i.e., B wins the bid)
- If B does not know C's true value, then $\max_x \min_y u(x, y) = 0$ since C could bid more than x_0 and $\min_y u(x, y) = 0$, so that there is no incentive for B to bid when using this principle
- If B knows that C's true value is $y_0 < x_0$, then the maximin solution for B is a value slightly larger than y_0

SEQUENTIAL GAMES

- In sequential games, the participants make decisions over time, usually in alternation
- The payoffs could accrue cumulatively during the sequence of play, as with tricks taken in the card game bridge, or the payoff may be determined only at the end of the sequence, as with the checkmate in a game of chess
- Often, the payoffs are stochastic
- We focus on two-person sequential games with perfect information, meaning that, at every stage, each opponent knows the choice that was made by the other

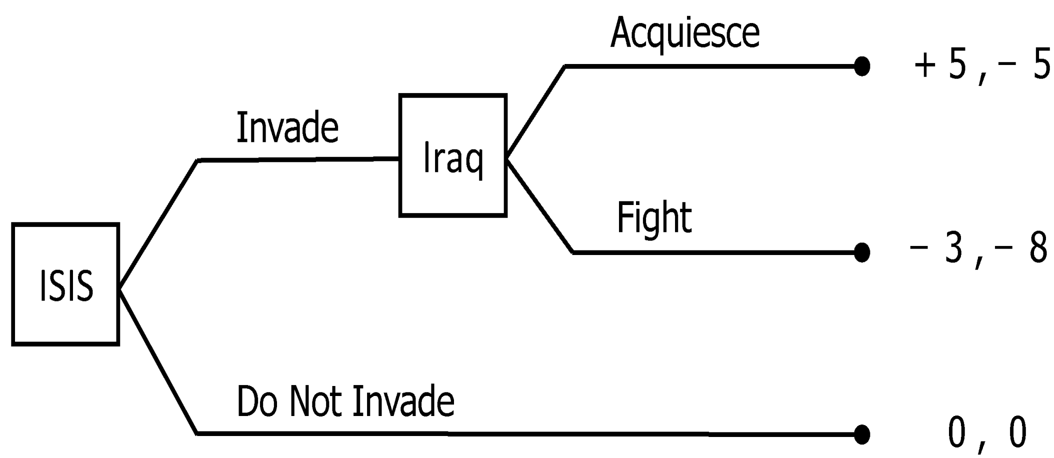

SEQUENTIAL GAMES

Sequential games are specified in terms of several elements:

- A set of n opponents (often $n = 2$)
- A game tree, which describes how the game develops, with decision nodes, chance nodes and terminal nodes
- Decision node labels, indicating which player owns each node
- A set of moves at each decision node, indicating the choices available to the node owner
- Distributions at each chance node, describing the values that may be taken by the corresponding random variable
- Payoffs at the end of each path in the game tree, determining the outcome of the game

SEQUENTIAL GAMES

- Suppose the Islamic State in Iraq and Syria (ISIS) is deciding whether or not to annex territory in Iraq
- If they invade, then Iraq must decide whether to fight or acquiesce
- This example assumes that the payoffs are non-random and known to both parties, so the game tree is actually a decision tree, since there are no chance nodes
- Payoffs in billions of dollars: first is ISIS's, and the second is Iraq's
- If ISIS invades and Iraq chooses to fight, then ISIS wins (a defeat is excluded!) and conquers the territory, worth \$5 billion, but both parties pay a cost for conflict (\$3 billion for Iraq and \$8 billion for ISIS)

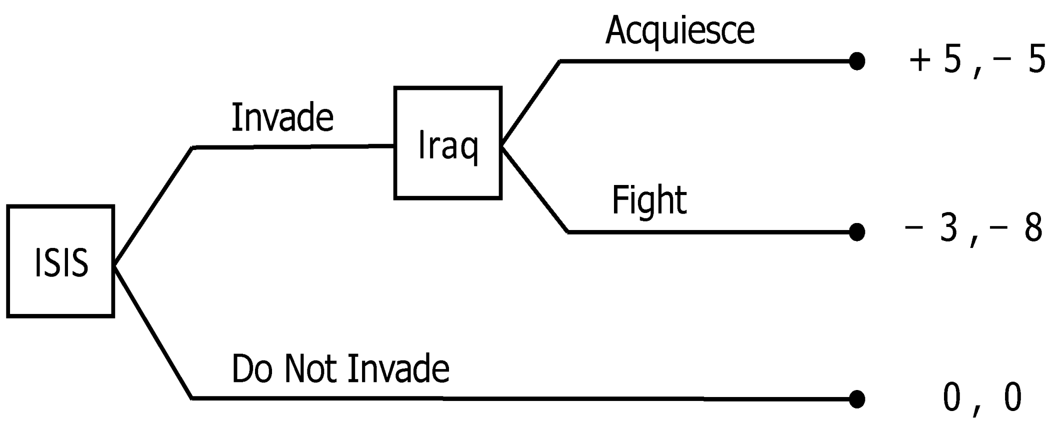

SEQUENTIAL GAMES

- This is a very simple sequential game where ISIS decides first and then Iraq decides later (if there is the invasion)

- In this case it is possible to present a bimatrix with the payoffs

- There are two Nash equilibria but one (Not Invade, Fight) is impossible while the other (Invade, Acquiesce) implies that ISIS will attack and Iraq will surrender (Iraq fighting back is worse for both sides!)

	Fight	Acquiesce
Invade	$(-3, -8)$	$(5, -5)$
Not Invade	$(0, 0)$	$(0, 0)$

SEQUENTIAL GAMES

- We now present a Defend-Attack-Defend game
- In this game the Defender (she) first deploys her defensive resources, then the Attacker (he) observes the deployment and chooses his attack and, finally, the Defender attempts to recover from the attack as best she can
- We present both an influence diagram and a game tree
  - Nodes D_1 and D_2 correspond to the Defender's first and second decisions, i.e., $d_1 \in \mathcal{D}_1$ and $d_2 \in \mathcal{D}_2$, respectively
  - Node A represents the Attacker's decision $a \in \mathcal{A}$
  - S is a shared uncertainty node about the success of the attack

SEQUENTIAL GAMES

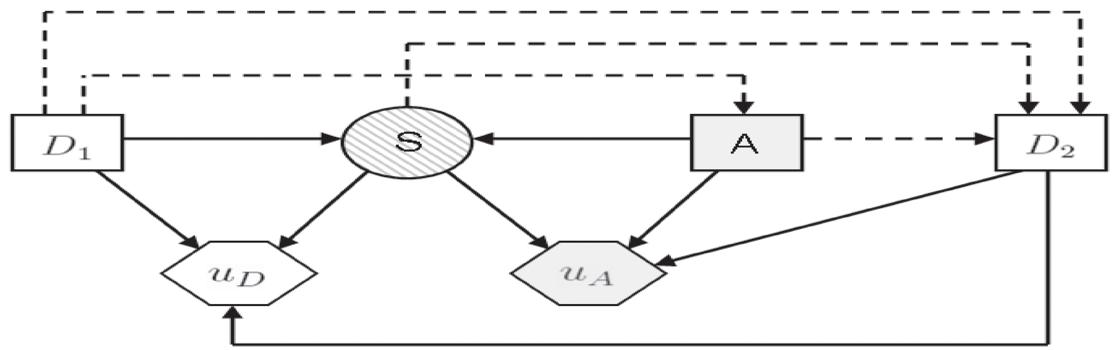

(a) Multi-agent influence diagram for the Defend-Attack-Defend game.

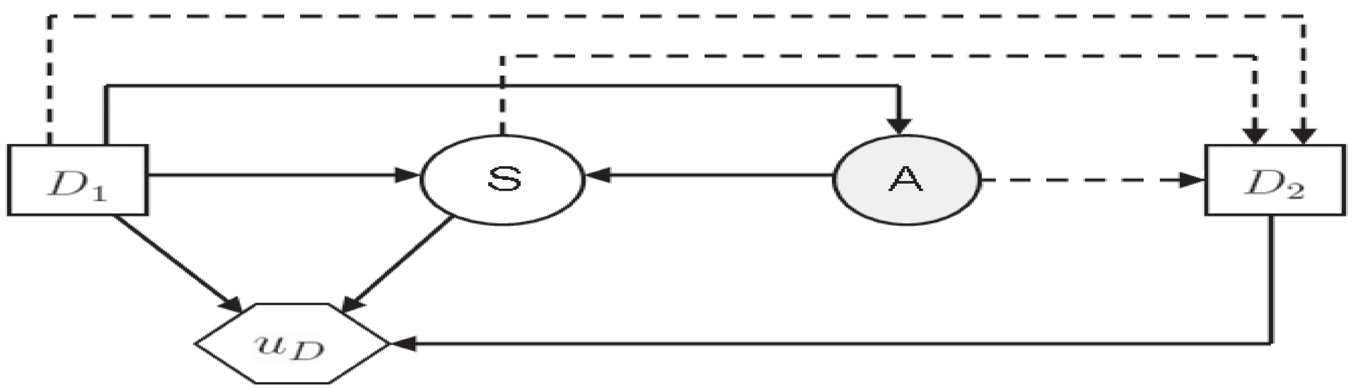

(b) The Defend-Attack-Defend game tree.

SEQUENTIAL GAMES

- In the previous slide the MAID shows the information available to each opponent at each stage of the game, and the relationships between utilities, decisions, and chance, while the game tree shows the sequence of play, which is helpful in backwards induction (the computational technique used to find the optimal decisions)

- As written, this model assumes that the only relevant uncertainty is the success level S of the attack, which depends probabilistically on $(d_1, a) \in \mathcal{D}_1 \times \mathcal{A}$

- The payoff to the Defender depends on (d_1, s, d_2): the cost of her initial defense, the success of the attack, and the success of the recovery effort

- The payoff to the Attacker depends on (a, s, d_2): the effort in mounting his attack, its success, and the success of her recovery effort

- The model may be easily generalised so that the outcome of the recovery effort is also a random variable

SEQUENTIAL GAMES

- Game theory needs the Defender to know the Attacker's utilities and probabilities, and the Attacker to know the Defender's utilities and probabilities, and for both to know that these are common knowledge

- Often, the utility functions and probability assessments depend upon all of d_1, s, a, and d_2 (i.e., the utility and probability are affected not only by the chance outcome, but also by the choices of each opponent)

- However here we make the simplifying assumption that utilities and probabilities depend only upon the outcome of the opponent's decision, but not also on the decision itself

- In such cases, denote the Defender's and Attacker's utility functions by $u_D(d_1, s, d_2)$ and $u_A(a, s, d_2)$, respectively, and their probability assessments about the success of attack by $p_D(S = s|d_1, a)$ and $p_A(S = s|d_1, a)$, respectively

SEQUENTIAL GAMES

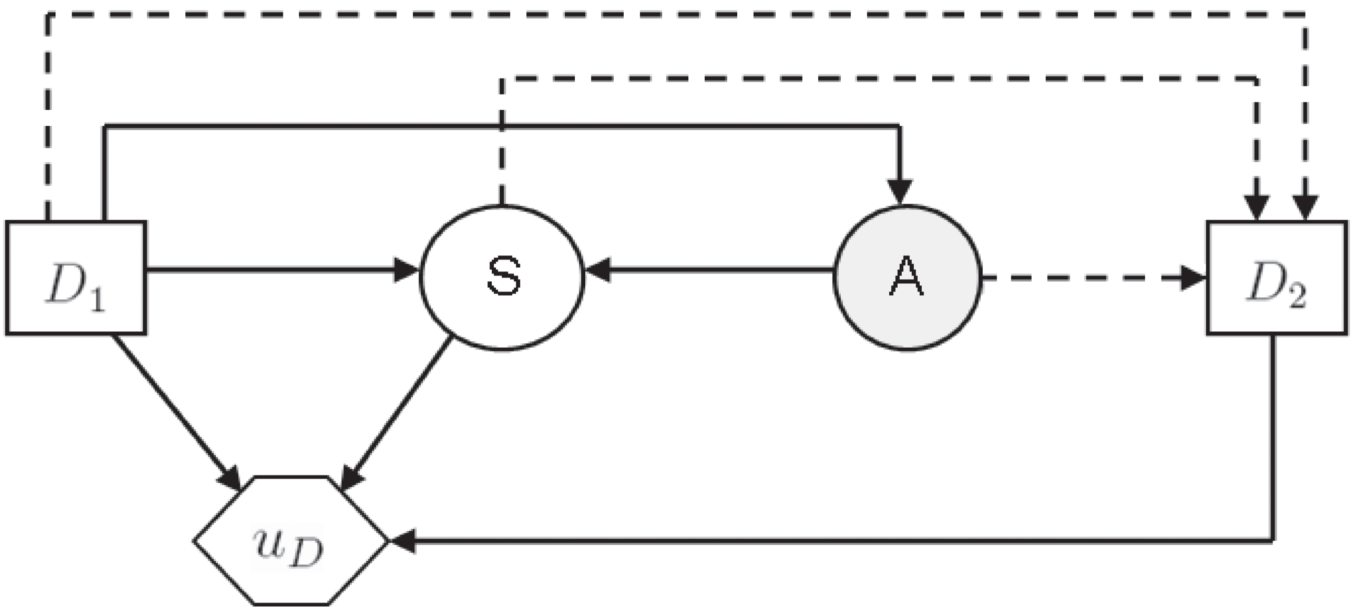

SEQUENTIAL GAMES

- Intuition would suggest to choose first the Defender's initial optimal defense d_1^* , then the optimal attack $a^*(d_1^*)$ after observing d_1^* and, finally, the optimal recovery action $d_2^*(d_1^*, s)$, after observing the consequence s of the previous actions

- But this is computationally unfeasible $\Rightarrow$ go backwards and use backwards induction
  - Find the optimal defense $d_2^*(d_1^*, s)$ for any pair (d_1, s)
  - For each pair (d_1, a), both Defender and Attacker can compute their expected utilities, since they have specified their own distributions of the consequences s for each (d_1, a) and their own utilities associated to their actions (d_1 for Defender and a for Attacker, consequence s and optimal decision d_2^*)
  - Both expected utilities are known to both opponents (common knowledge!)
  - For each initial decision d_1, the Attacker finds his optimal attack $a^*(d_1)$
  - Finally, the Defender chooses her optimal initial defense d_1^*, knowing which optimal attack will be chosen by the opponent and her optimal recovery action

SEQUENTIAL GAMES

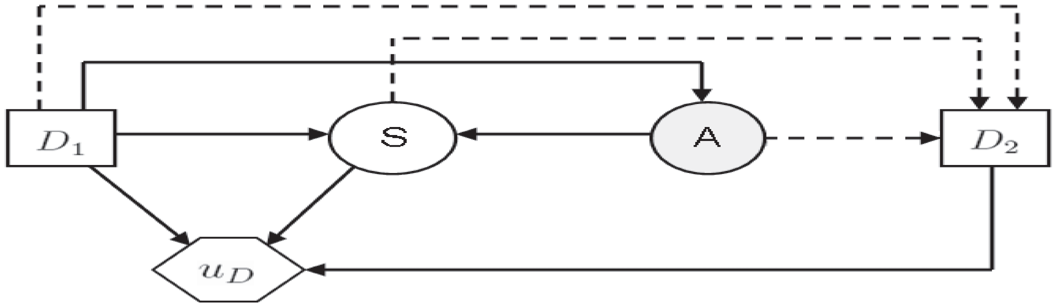

- Using backwards induction, at node D_2 for the game tree above, the Defender's best response after observing $(d_1, s) \in \mathcal{D}_1 \times \mathcal{S}$ is $d_2^*(d_1, s) = \arg\max_{d_2 \in \mathcal{D}_2} u_D(d_1, s, d_2)$

- Under the common knowledge assumption, the Defender's behaviour at D_2 will be anticipated by the Attacker

- At node S, Defender's expected utility associated with each $(d_1, a) \in \mathcal{D}_1 \times \mathcal{A}$: $\psi_D(d_1, a) = \int u_D(d_1, s, d_2^*(d_1, s)) p_D(S = s | d_1, a) ds$

- At node S, Attacker's expected utility associated with each $(d_1, a) \in \mathcal{D}_1 \times \mathcal{A}$: $\psi_A(d_1, a) = \int u_A(a, s, d_2^*(d_1, s)) p_A(S = s | d_1, a) ds$

- Both expected utilities are known to both opponents

SEQUENTIAL GAMES

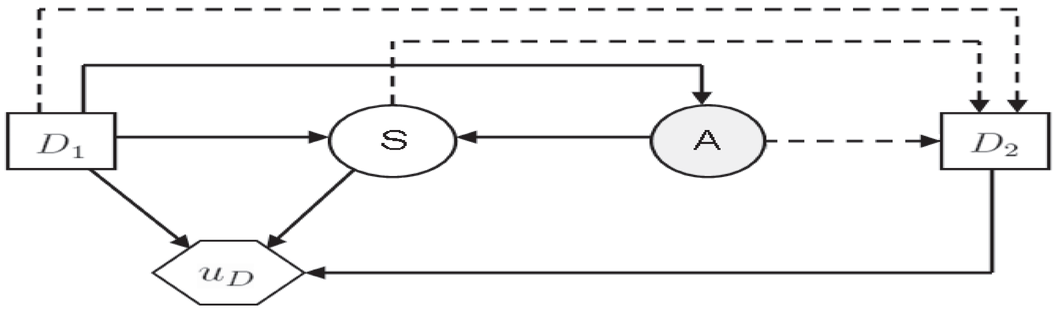

- The Attacker now finds his optimal attack at node A, after observing the Defender's move $d_1 \in \mathcal{D}_1$, by solving $a^*(d_1) = \arg\max_{a \in \mathcal{A}} \psi_A(d_1, a)$
- Knowing this, the Defender finds her maximum expected utility decision at node D_1 through $d_1^* = \arg\max_{d_1 \in \mathcal{D}_1} \psi_D(d_1, a^*(d_1))$, which gives the solution to the game
- Assuming common knowledge, rational players, and perfect information, Game Theory prescribes that
  - The Defender should choose $d_1^* \in \mathcal{D}_1$ at node D_1
  - The Attacker should choose attack $a^*(d_1^*) \in \mathcal{A}$ at node A after observing d_1^*
  - The Defender, after observing $s \in \mathcal{S}$, should choose $d_2^*(d_1^*, s) \in \mathcal{D}_2$ at node D_2

ARA FOR SEQUENTIAL GAMES

- ARA can be used in sequential games, and it allows one to drop the assumption that all utilities and probabilities are common knowledge

- From the ARA perspective, in the Defend-Attack-Defend model, the Attacker's decision at node A is uncertain to the Defender, and she must model her uncertainty through a random variable

- This is reflected in the influence diagram and the game tree (next slide), where the Attacker's decision node A has been converted to a chance node

- The Defender needs to assess $p_D(a|d_1)$, her predictive distribution about which attack a will be chosen at node A for each $d_1 \in \mathcal{D}_1$

- Additionally, she must make the (more standard) assessments about $u_D(d_1, s, d_2)$ and $p_D(s|d_1, a)$

ARA FOR SEQUENTIAL GAMES

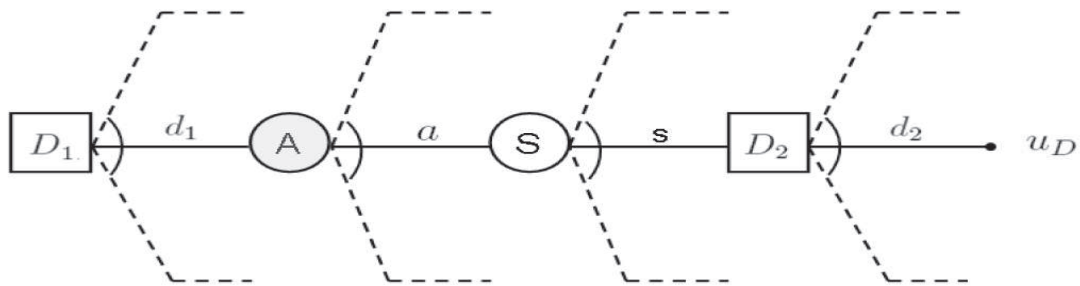

(a) The influence diagram for the ARA solution of the Defend-Attack-Defend game.

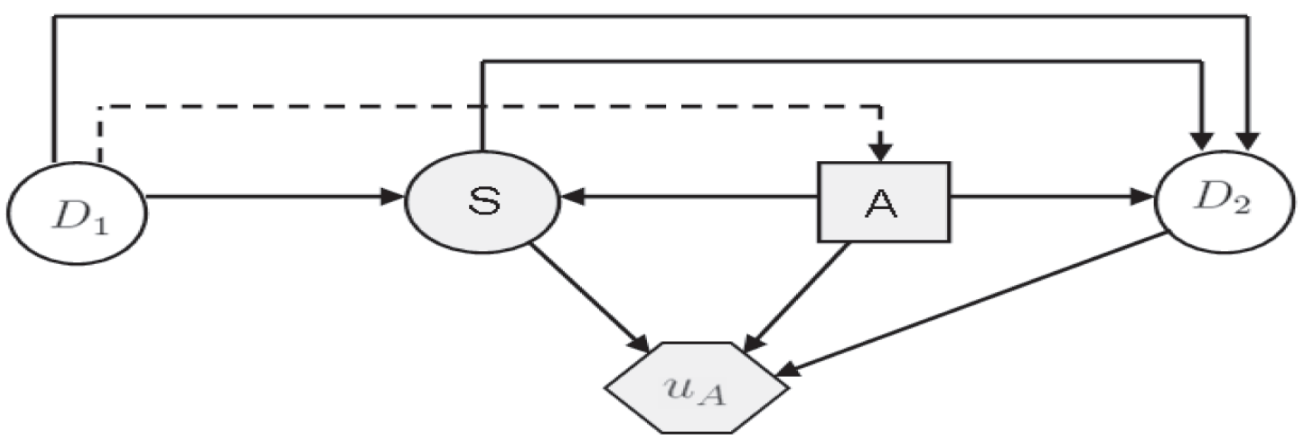

(b) The game tree for the ARA solution of the Defend-Attack-Defend game.

ARA FOR SEQUENTIAL GAMES

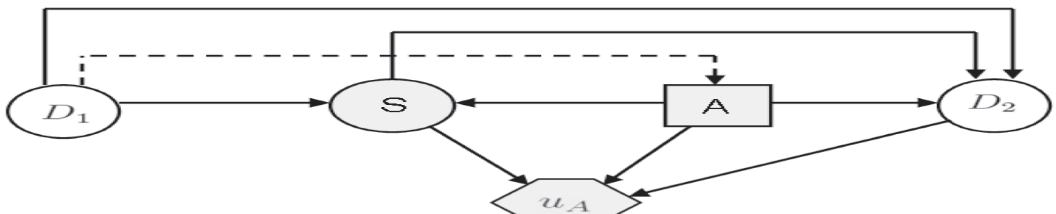

- The Defender works backward in the tree to solve her problem

- At node D_2 she computes her optimal action $d_2^*(d_1, s)$ for each $(d_1, s) \in \mathcal{D}_1 \times \mathcal{S}$, with associated utility $\psi_D(d_1, a) = \int u_D(d_1, s, d_2^*(d_1, s))p_D(S = s|d_1, a)ds$

- Then, at node S, for each $(d_1, a) \in \mathcal{D}_1 \times \mathcal{A}$, she finds her expected utility: $\psi_D(d_1, a) = \int u_D(d_1, s, d_2^*(d_1, s))p_D(S = s|d_1, a)ds$, for each $(d_1, a) \in \mathcal{D}_1 \times \mathcal{A}$

- Then, her subjective probability assessment of what the Attacker will do, $p_D(A|d_1)$, is used to compute her expected utility at node A for each $d_1 \in \mathcal{D}_1$: $\psi_D(d_1) = \int \psi_D(d_1, a)p_D(A = a|d_1)da$

- The Defender finds her maximum expected utility at D_1 as $d_1^* = \arg\max_{d_1 \in \mathcal{D}_1} \psi_D(d_1)$

- Thus, backwards induction shows that the Defender's best strategy is to choose first d_1^* at node D_1, and later, after observing $s \in \mathcal{S}$, choose $d_2^*(d_1^*, s)$ at node D_2

ARA FOR SEQUENTIAL GAMES

- As usual, ARA requires the decision maker to assess $p_D(a|d_1)$. This could be done through some form of risk analysis (non-strategic behaviour), or by developing a model for the strategic analysis of the opponent

- Given $p_D(a|d_1)$, the analysis of the Attacker's decision problem, as seen by the Defender, is shown in the next two figures, where the Attacker's probabilities and utilities are assessed by the Defender

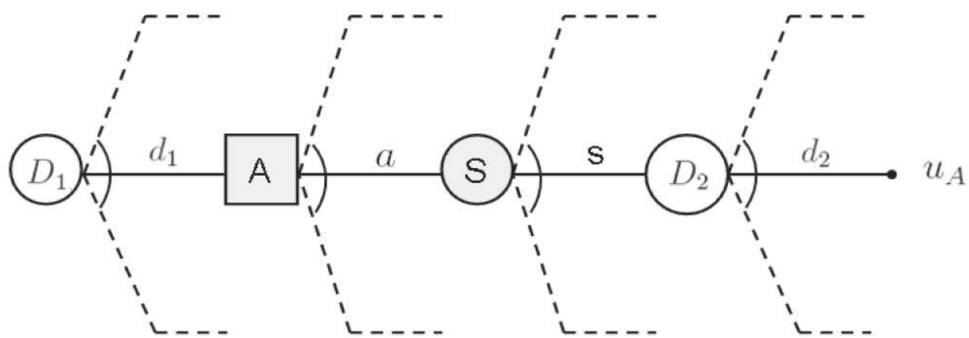

(a) Influence diagram for the Defender's view of the Attacker's problem.

ARA FOR SEQUENTIAL GAMES

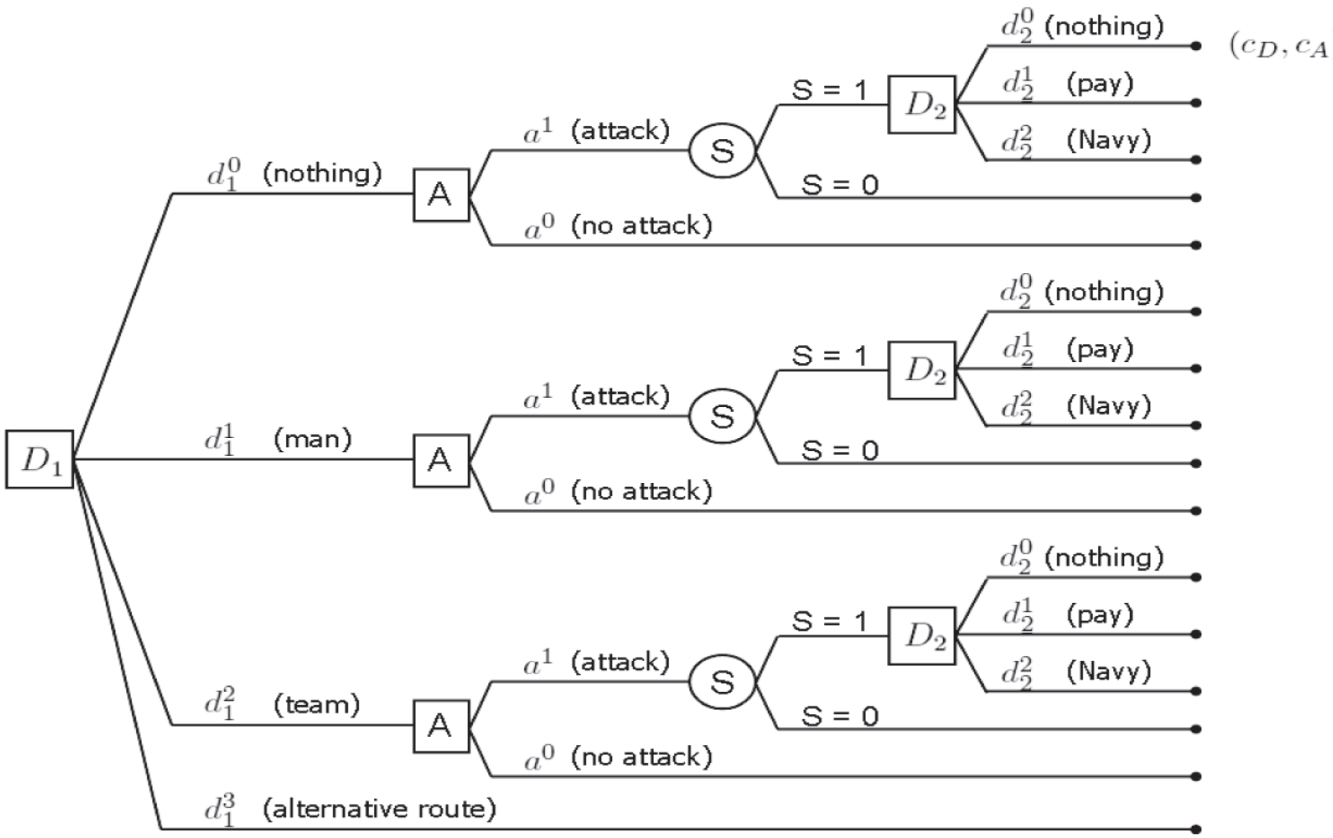

(b) Game tree for the Defender's view of the Attacker's problem.

ARA FOR SEQUENTIAL GAMES

- The assessment of the Attacker's probabilities and utilities may be based upon historical data and expert opinion, or, if those salient information is not available, then the Defender may choose to use a noninformative distribution $p_D(a|d_1)$

- In the former case, in order to elicit $p_D(a|d_1)$, the Defender must assess $u_A(a, s, d_2)$, $p_A(s|d_1, a)$, and $p_A(d_2|d_1, a, s)$, i.e., Attacker's probabilities and utilities

- In general, she does not know these quantities, and represents her uncertainty through a joint distribution F on the random quantities $U_A(a, s, d_2)$, $P_A(s|d_1, a)$, and $P_A(d_2|d_1, a, s)$

- These distributions might be elicited in various ways, e.g. Dirichlet process centered at Defender's probabilities

- The Defender then solves her perception of the Attacker's decision problem using backwards induction over the game tree, propagating her uncertainty, encoded by F, to get distributions over the random action $A^*(d_1)$ for each d_1

ARA FOR SEQUENTIAL GAMES

Specifically, if all choice sets are continuous, the Defender solves as follows

- At decision node D_2, compute the random expected utilities
$(d_1, a, s) \rightarrow \Psi_A(d_1, a, s) = \int U_A(a, s, d_2) P_A(D_2 = d_2 | d_1, a, s) dd_2$

- At chance node S, compute the random expected utilities
$(d_1, a) \rightarrow \Psi_A(d_1, a) = \int \Psi_A(d_1, a, s) P_A(S = s | d_1, a) ds$

- At decision node A , compute the random optimal initial decision
$d_1 \rightarrow A^*(d_1) = \arg\max_{a \in \mathcal{A}} \Psi_A(d_1, a)$

- Thus, the Defender's predictive distribution over attacks, conditional on her first defensive decision d_1, is given by $\int_0^a p_D(A = x | d_1) dx = Pr[A^*(d_1) \leq a]$

- If the choice sets are not continuous, the Defender would reason similarly, but replacing integrals with sums and finding the predictive distribution over a set

ARA FOR SEQUENTIAL GAMES

- When the beliefs are complex, the Defender's predictive distribution can be approximated using Monte Carlo simulation

- Specifically, for each $d_1 \in \mathcal{D}_1$, the following algorithm finds the Defender's belief about the probability of each attack

Do, for $i = 1, \ldots, N$:

Draw $\big(U_A^i(a,s,d_2), P_A^i(S \mid d_1,a), P_A^i(D_2 \mid d_1,a,s)\big) \sim F$.

At chance node D_2, compute

$$(d_1,a,s) \to \psi_A^i(d_1,a,s) = \int U_A^i(a,s,d_2)\, P_A^i(D_2 = d_2 \mid d_1,a,s)\, dd_2.$$

At chance node S, compute

$$(d_1,a) \to \psi_A^i(d_1,a) = \int \psi_A^i(d_1,a,s)\, P_A^i(S = s \mid d_1,a)\, ds.$$

At decision node A, compute

$$d_1 \to a_i^*(d_1) = \operatorname{argmax}_{a \in \mathcal{A}} \psi_A^i(d_1,a).$$

Approximate $\int_0^a p_D(A = x \mid d_1) dx$ by $\#\{a_i^*(d_1) \ : \ a_i^*(d_1) \leq a\}/N$
where $\#\{\cdot\}$ is the cardinality of the set.

ARA FOR SEQUENTIAL GAMES

- Given F, the Defender's assessment of $p_D(a|d_1)$ is straightforward

- In many situations, it is relatively simple to elicit the random utility $U_A(a, s, d_2)$ and the random probability $P_A(s|d_1, a)$ in F

- However, the assessment of $P_A(d_2|d_1, a, s)$ within F can be problematic; the Defender should exploit any information she has on how the Attacker analyses her decision problem

- She may think he seeks a Nash equilibrium, or uses level-k thinking; if she has little insight, she can represent that through some high-variance or non-informative distribution

CASE STUDY: SOMALI PIRATES

- As a concrete example of sequential ARA, consider the decisions made in defending a ship from piracy
- Between 2005 and 2011, pirates threatened shipping in the Gulf of Aden
- No ship within several hundred miles of the Somali coast was safe and it was a major issue in international security
- Economic and political changes have led many Somali fishermen to take up piracy, and infrastructure evolved to support this new enterprise
- The infrastructure had multiple agents: village elders who act as the de facto local government, Somali businessmen who invest in piracy, and negotiators who broker ransoms for captured ships and crews
- So an analyst can model the pirates as pragmatic businessmen who pursue their goals strategically

CASE STUDY: SOMALI PIRATES

- A typical attack is undertaken by a small groups of about ten men in fast boats which depart from a mother-ship
- If successful, about 50 pirates are left on-board to pilot the captured ship into harbor, while another 50 or so pirates provide logistical support from the shore
- The goal is ransom rather than theft; it is more profitable, and the pirates reinvest part of their gains in equipment and training
- First consider anti-piracy strategy from the perspective of a ship owner, structuring the problem as a game tree
- This leads to a sequential Defend-Attack-Defend game between the Owner and the Pirates, following the formulation in Sevillano, Rios Insua and Rios (2012).

CASE STUDY: SOMALI PIRATES

- For the first defensive decision, the Owner may choose among many options, including various levels of on-board armed security and selecting an alternate (longer) route that avoids the Gulf of Aden

- Next, the Pirates, who have a network of spies that provide information about security, cargo and crew, respond to the Owner's initial decision by either attacking or not attacking the ship

- If the pirate attack is successful, the Owner will have to decide how much to pay in ransom, or perhaps she will hire armed forces to re-seize the ship

- Specifically, assume that the Owner can select one of the following four defensive actions (i.e., elements of $\mathcal{D}_1$):
  - d_1^0: Do nothing, i.e., no defensive action is taken
  - d_1^1: Hire an armed guard to travel on the ship
  - d_1^2: Hire a team of two armed guards for the ship
  - d_1^3: Use the Cape of Good Hope route, not the Suez Canal

CASE STUDY: SOMALI PIRATES

- Once the Owner has made her initial choice, the Pirates observe it and decide whether or not to attack (a^1 and a^0, respectively, in $\mathcal{A}$)

- An attack either results in the ship being hijacked ($S = 1$) or not ($S = 0$), with probabilities that depend on the Owner's initial choice

- If the ship is hijacked, then the Owner has three possible responses (elements of $\mathcal{D}_2$):
  - d_2^0: Do nothing, i.e., refuse to pay the Pirates' ransom
  - d_2^1: Pay a ransom, thus recovering the ship and crew
  - d_2^2: Pay the Navy (i.e., military units) to recapture the ship and crew

- Obviously, this framing of the decision problem is only an approximation to the true complexity of the application, but it captures the salient elements

CASE STUDY: SOMALI PIRATES

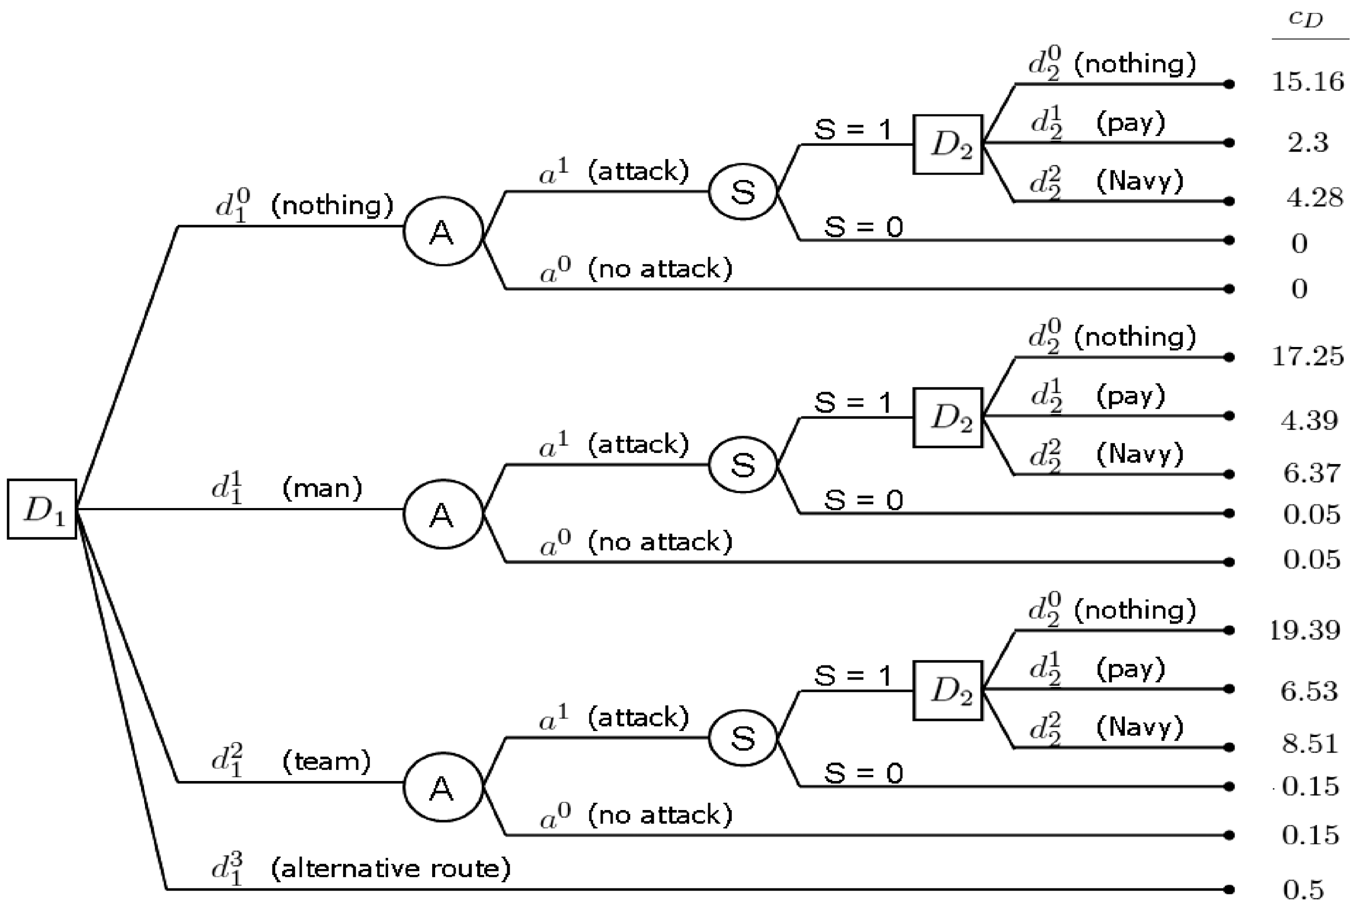

CASE STUDY: SOMALI PIRATES

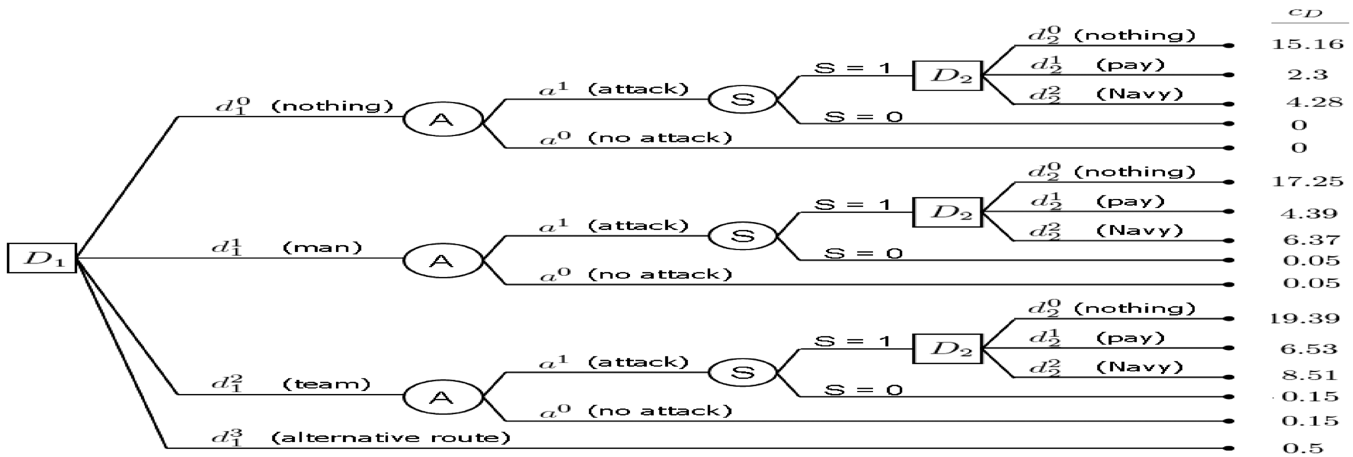

- The game tree represents the sequence of decisions and outcomes faced by the Owner and the Pirates

- The nodes D_1 and D_2 correspond to the Owner's first and second decisions respectively, the node A represents the Pirates' decision, and the chance node S represents the outcome of the attack (if undertaken)

- The pair (c_D, c_A) represents the consequences to the Owner and the Pirates, respectively, from the corresponding sequence of decisions and outcomes

CASE STUDY: SOMALI PIRATES

- The Owner's sequential ARA decision problem is mapped as a game tree in which the Pirates' decision node A has been replaced by a chance node

- This shows that the Pirates' decision is unknown in advance to the Owner, forcing the bulk of her modeling to focus on assessments of probabilities over the Pirates' actions conditional on her own

- Thus, to solve her decision problem, she needs to assess $p_D(a|d_1)$, her predictive probability of an attack given each $d_1 \in \mathcal{D}_1$

- Additionally, she needs to make assessments $p_D(s|d_1, a^1)$ and $u_D(c_D)$, with c_D representing her monetary cost equivalent from the multi-attribute consequences associated with each branch in the tree

- These assessments are easier to make because they are non-strategic, and also because in this application there is relevant historical information

CASE STUDY: SOMALI PIRATES

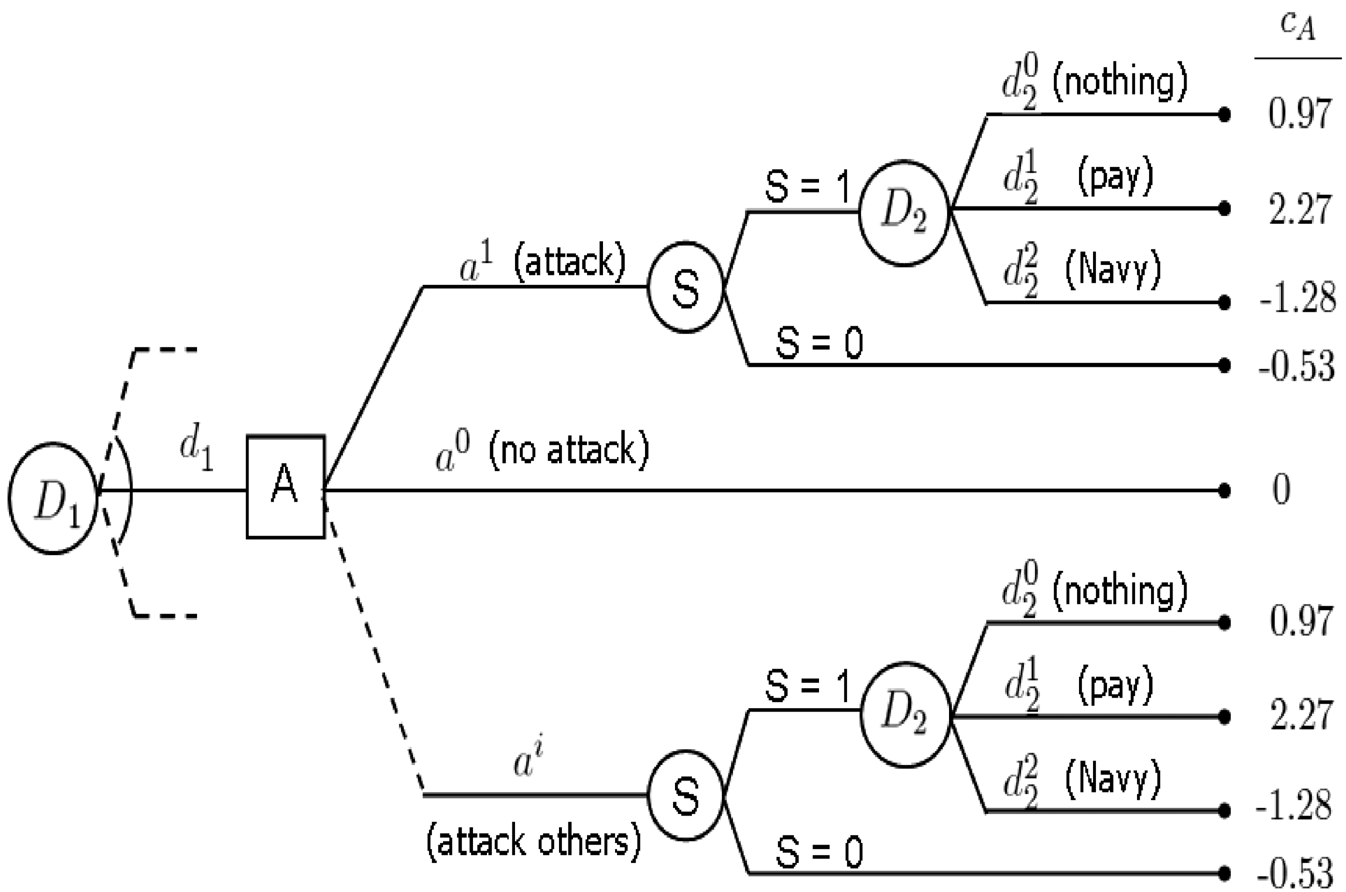

CASE STUDY: SOMALI PIRATES

- For the Owner, the possible consequences are loss and ransom of the ship, the costs associated with defense preparation, the costs associated with a possible battle, and the profit from a voyage that is either not attacked or which successfully repels armed boarders

- These costs are not all commensurate since, e.g., an attack, successful or unsuccessful, could include loss of life on either side

- This analysis assumes that the Owner puts no value on the life of a Somali pirate, and uses the value of a statistical life for each member of the crew (including any guards that may be hired)

- Note that there are many methods for estimating the value of a statistical life (e.g., the total income over the remainder of an expected lifespan)

- Now we provide Owner's direct costs for the defensive actions in $\mathcal{D}_1$ and $\mathcal{D}_2$

CASE STUDY: SOMALI PIRATES

- Owner's direct costs for the defensive actions in $\mathcal{D}_1$
  - €0, if the Owner hires no guards (i.e., d_1^0)
  - €0.05M, if the Owner hires one guard (i.e., d_1^1) and this includes the six-month salary, equipment, travel expenses, and so forth
  - €0.15M, if the Owner hires two guards (i.e., d_1^2) this includes the six-month salaries for two armed guards, with better equipment and associated expenses
  - €0.5M, if the Owner chooses to circumnavigate the Cape of Good Hope (i.e., d_1^3), since the increased distance entails higher operating expenses
- Owner's direct costs for the defensive actions in $\mathcal{D}_2$
  - €0, if the ship is attacked and the Owner pays no ransom (i.e., d_2^0). (This is the direct cost; the loss of the ship and cargo is accounted for separately later)
  - €2.3M, if the ship is successfully attacked and the Owner pays a ransom (i.e., d_2^1)- The average of some actual ransoms is €2.3M, and here the ransom to be paid, for simplicity, is assumed to be fixed although it could be treated as a r.v.
  - €0.2M, if the ship is attacked and the Owner calls for the Navy (i.e., d_2^2), as based on the fees for the international coalition ships already in the area

CASE STUDY: SOMALI PIRATES

- The Owner's direct cost for the deaths depends only on insurance and litigation but the ARA should take deaths directly into account

- For this example, the analysis assumes that if an attack is successfully repelled $(S = 0)$, then no lives are lost

- In a successful attack $(S = 1)$, the analysis assumes that all armed guards are killed and, depending on the chosen response at the node D_2 , there may be additional fatalities:

  – if the Owner does not ransom the ship, the angry Pirates kill four of the crew

  – if the Owner pays the ransom, no one else dies

  – if the hijacked ship is rescued by the Navy, there are two more deaths

- A slightly more complex analysis would properly treat the number of deaths and their costs as random variables, but this discussion omits that and follows Martinez and Mendez (2009) in fixing the statistical value of a (Spanish) life at €2.04M

- Similarly, we assume the depreciated value of the ship and its cargo is €7M

CASE STUDY: SOMALI PIRATES

e 4.2 The Owner's costs associated with different tree paths.

D_1	S	D_2	Ship loss	Action costs	Lives lost	c_D
do nothing	$S = 1$	don't pay	1	0 + 0	0 + 4	15.16
do nothing	$S = 1$	pay ransom	0	0 + 2.3M	0 + 0	2.30
do nothing	$S = 1$	call Navy	0	0 + 0.2M	0 + 2	4.28
do nothing	$S = 0$		0	0	0	0.00
hire guard	$S = 1$	don't pay	1	0.05M + 0	1 + 4	17.25
hire guard	$S = 1$	pay ransom	0	0.05M + 2.3M	1 + 0	4.39
hire guare	$S = 1$	call Navy	0	0.05M + 0.2M	1 + 2	6.37
hire guard	$S = 0$		0	0.05M	0	0.05
hire team	$S = 1$	don't pay	1	0.15M + 0	2 + 4	19.39
hire team	$S = 1$	pay ransom	0	0.15M + 2.3M	2 + 0	6.53
hire team	$S = 1$	call Navy	0	0.15M + 0.2M	2 + 2	8.51
hire team	$S = 0$		0	0.15M	0	0.15
d_1^3 (alternative route)			0	0.5M	0	0.50

CASE STUDY: SOMALI PIRATES

- The owner is risk averse, with utility function $u_D(c_D) = 1 - \exp(-\alpha c_D)$, $\alpha > 0$
- For sensitivity analysis ARA considers optimal decisions for $\alpha \in \{0.1, 0.4, 1, 2, 5\}$
- Based on historical information, the ARA assumes that the Owner believes $p_D(S = 1|a^1, d_1^0) = 0.4$. i.e. a successful attack when no armed guards are hired
- It also assumes a successful attack with $p_D(S = 1|a^1, d_1^1) = 0.1$ with one guard, and $p_D(S = 1|a^1, d_1^2) = 0.05$ with two guards
- In order to implement an ARA, the Owner must estimate the probability of attack conditional on each of her initial defensive choices
- Using historical data, the probability of an attack is about 0.005 but this does not account for the defensive choices, nor the intelligence network used by the Pirates to identify ships with valuable cargo and vulnerabilities, so using 0.005 as the attack probability implies a non-strategic opponent

CASE STUDY: SOMALI PIRATES

- We assume that the Owner performs a level-2 analysis
- Assume that the Owner believes that the Pirates are expected utility maximisers who derive the Owner's uncertainty about the Pirates' decision on whether to attack from her uncertainty about the Pirates' probabilities and utilities
- Thus, the Owner must analyse the decision problem from the Pirates' perspective

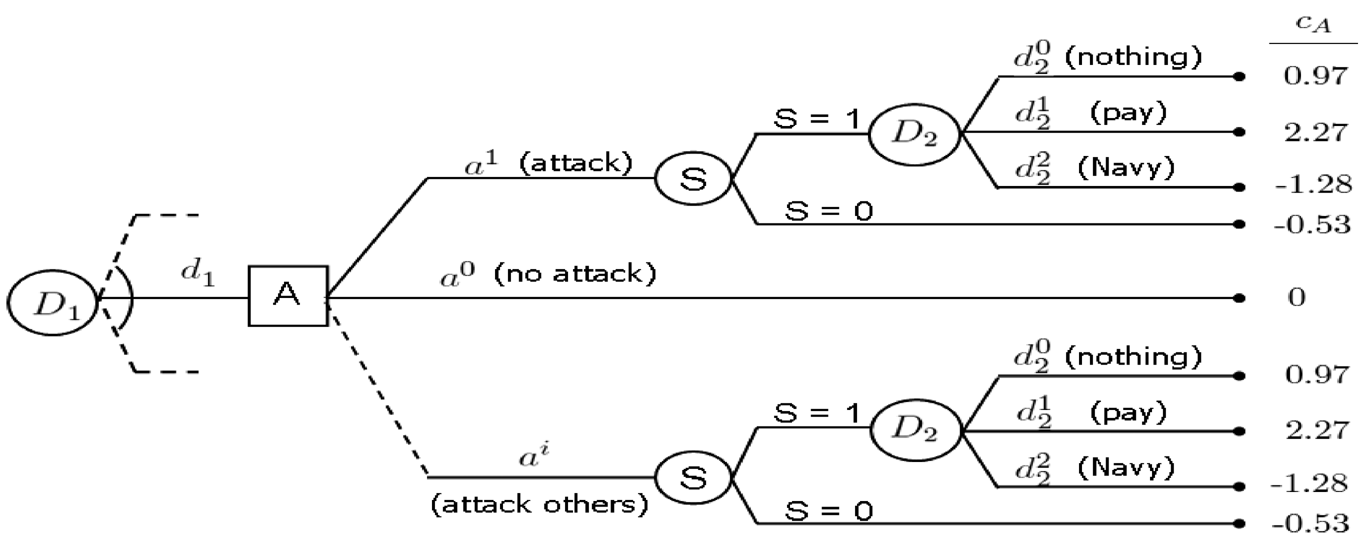

CASE STUDY: SOMALI PIRATES

- The set of alternatives for the Pirates has expanded to include alternatives $a^i \in \mathcal{A}$, for $i = 2, \ldots, n$, representing the Pirates' option to attack ships owned by others
- Also, there are new chance nodes D_2 at the end of the tree paths starting at a^i, which represent the response of ships $i = 2, \ldots, n$ to a hijacking attempt, since these are uncertainties from the Pirates' standpoint
- The Owner's analysis of the Pirates' decision making enables probabilistic assessment of the perceived preferences of the Pirates
- The Owner's uncertainty over these preferences is modeled through the random utility function $U_A(a, s, d_2)$, for $a \in \mathcal{A} = \{a^0, a^1, \ldots, a^n\}$
- The Owner's uncertainty about the Pirates' beliefs regarding a successful attack on her ship and the subsequent payoff is modeled by the random variables $P_A(S = 1|a^1, d^1)$ and $P_A(D_2|d_1, a^1, S = 1)$
- Similarly, The Owner's uncertainty about the Pirates' beliefs regarding attacks upon other ships is modeled by $P_A(S = 1|a^i)$ and $P_A(D_2|a^i, S = 1)$ for $i = 2, \ldots, n$
- To begin, we assume that the n ships are similar

CASE STUDY: SOMALI PIRATES

- The Owner thinks that the relevant outcomes for the Pirates are the net assets gained and the number of lives lost
- Carney (2009) reports that the average cost of an attack is about €30,000 and the average ransom paid is about €2.3M
- The Owner assumes that two pirates are killed when an attack is repelled, no pirates die in a successful attack and five pirates die if the Owner calls in the Navy
- We could treat these outcomes as random variables, through a slightly more elaborate analysis
- The Owner knows that her ship is less valuable to the Pirates than it is to her
- The Pirates can sell its cargo, sell it for scrap, or use it as a mother ship for new attacks
- So the Owner assesses the economic value of the ship to the Pirates at €1M
- Also, suppose the Owner thinks that Pirates value the life of one of their own at €0.25M

CASE STUDY: SOMALI PIRATES

Table 4.3 Payoffs to the Pirates according to each paths in their game tree, $i = 1, \ldots, n$. These calculations assume that all ships are equivalent in expectation for all relevant risks and rewards.

A	S	D_2	Ship kept	Profit	Lives lost	c_A
no attack			0	0	0	0.00
attack	$S=1$	don't pay	1	−0.03M	0	0.97
attack	$S=1$	pay ransom	0	2.27M	0	2.27
attack	$S=1$	call Navy	0	−0.03M	5	−1.28
attack	$S=0$		0	−0.03M	2	−0.53

- The aggregated monetary equivalent is c_A, in the last column

CASE STUDY: SOMALI PIRATES

- The Owner also needs to model her beliefs about the extent to which Pirates are risk-seeking with respect to profits

- She thinks they have utility function of the form $u_A(c_A) = \exp(c \times c_A)$, with $c > 0$

- The Owner is not sure about c and chooses a uniform distribution on $[0, 20]$ for it

- The uncertainty over c induces uncertainty over u_A, providing U_A

- The Owner uses subjective distributions to model Pirates' beliefs about the probability of a successful attack, conditional on her observed defensive choices

- Suppose she assumes that $P_A(S = 1|a^1, d_1^0) \sim \mathcal{B}eta(40, 60)$ with no defensive action taken, that $P_A(S = 1|a^1, d_1^1) \sim \mathcal{B}eta(10, 90)$ when one armed guard is hired, and that $P_A(S = 1|a^1, d_1^2) \sim \mathcal{B}eta(50, 950)$ when two guards are hired

- Here, the expected values of these distributions equal the assessed probabilistic beliefs of the Defender for the same situation, reflecting their common knowledge of past exploits, but in other situations the means could be different

CASE STUDY: SOMALI PIRATES

- For the other ships, the Owner does not know whether they carry armed guards, nor, if they do, whether there are one or two of them

- So, to her, the probability that the Pirates assign to a successful attack on another ship is a mixture of the three previous beta distributions and, in absence of more precise information, she puts equal weight on all three components

- Finally, for the level-2 ARA, the Owner must assess how the Pirates think she will respond to a successful attack

- It is reasonable to imagine that the Pirates believe her initial decision is a clue to her response

- If her first decision were aggressive (i.e., to hire two armed guards), then her response to an attack is also likely to be aggressive (i.e., to call in the Navy)

CASE STUDY: SOMALI PIRATES

- To reflect this concretely, the Owner believes that the Pirates have Dirichlet distributions over her three response options (write off the loss, pay the ransom, or call in the Navy)
  - When she did not hire guards, she thinks the Pirates believe that all three responses are equally likely and use $P_A(D_2|d_1^0, A = a^1, S = 1) \sim \mathcal{D}ir(1, 1, 1)$
  - When she hired one guard, she thinks the Pirates believe that she is more likely to call in the Navy and use $P_A(D_2|d_1^1, A = a^1, S = 1) \sim \mathcal{D}ir(0.1, 4, 6)$
  - If she hired two guards, then $P_A(D_2|d_1^2, A = a^1, S = 1) \sim \mathcal{D}ir(0.1, 1, 10)$, implying she is much more likely to call in the Navy

- As before, for the other ships that are possible targets, the Owner places a mixture of Dirichlet distributions over the Pirates' beliefs about their responses and it is convenient to assume that $P_A(D_2|A = a^i, S = 1)$ puts equal weight on all three components

CASE STUDY: SOMALI PIRATES

- With all this machinery, it is now possible for the Owner to find the level-1 solution to the Pirates' decision problem, using backwards induction to solve her perception of the problem

1. She computes the random expected utilities corresponding to the Pirates' selection of a^1, conditional on her initial defense choices $d_1 \in \mathcal{D}_1 \setminus \{d_1^3\}$:

$$\Psi_A(d_1, a^1) = \left[\sum_{d_2 \in \mathcal{D}_2} U_A(a^1, S=1, d_2) \; P_A(D_2 = d_2 \mid d_1, a^1, S=1) \right] \times$$
$$P_A(S=1 \mid d_1, a^1) + P_A(S=0 \mid d_1, a^1) \; U_A(a^1, S=0).$$

2. She computes the random expected utilities corresponding to the Pirates' selection of a^i for $i = 2, \ldots, n$:

$$\Psi_A(a^i) = \left[\sum_{d_2 \in \mathcal{D}_2} U_A(a^i, S=1, d_2) \; P_A(D_2 = d_2 \mid a^i, S=1) \right] \times$$
$$P_A(S=1 \mid a^i) \; + \; P_A(S=0 \mid a^i) \; U_A(a^i, S=0).$$

3. She computes the Owner's predictive probabilities of being attacked ($A = a^1$) conditional on the initial defense choices, $d_1 \in \mathcal{D}_1 \setminus \{d_1^3\}$:

$$p_D(A = a^1 \mid d_1) = \Pr[\, \Psi_A(d_1, a^1) > \max\{U_A(a^0), \; \Psi_A(a^2), \ldots, \Psi_A(a^n)\} \,].$$

CASE STUDY: SOMALI PIRATES

- The probabilities are approximated through Monte Carlo simulation by drawing a sample $\{(u_A^k, p_A^k)\}_{k=1}^N \sim (U_A, P_A)$ from the Pirates' random utilities and probabilities, as assessed by the Owner, and solving the Pirates' decision problem for each draw

- This gives an estimate of the probability that the Owner's ship is chosen as the target

- Thus, with n possible ships to attack and N Monte Carlo draws, the estimate is
$$\frac{\#\{1 \leq k \leq N : \Psi_A^k(d_1, a^1) > \max\{u_A^k(a^0), \Psi_A^k(a^2), \ldots, \Psi_A^k(a^n)\}\}}{N}$$

- To illustrate, suppose there are nine possible ships that could be attacked

- Under the modeling assumptions described, $50,000$ Monte Carlo draws finds that $\widehat{p}_D(A = a^1|d_1^0) = 0.303$ is the estimated probability that the Owner's ship will be attacked if she does not hire guards, $\widehat{p}_D(A = a^1|d_1^1) = 0.026$ if she hires one armed guard, and $\widehat{p}_D(A = a^1|d_1^2) = 0.00004$ if she hires two armed guards

- So the Owner's estimate of the probability of attack decreases rapidly when she obtains protection

CASE STUDY: SOMALI PIRATES

- Given the estimated values of her beliefs about the attack probabilities, she solves her problem using backwards induction on the tree below

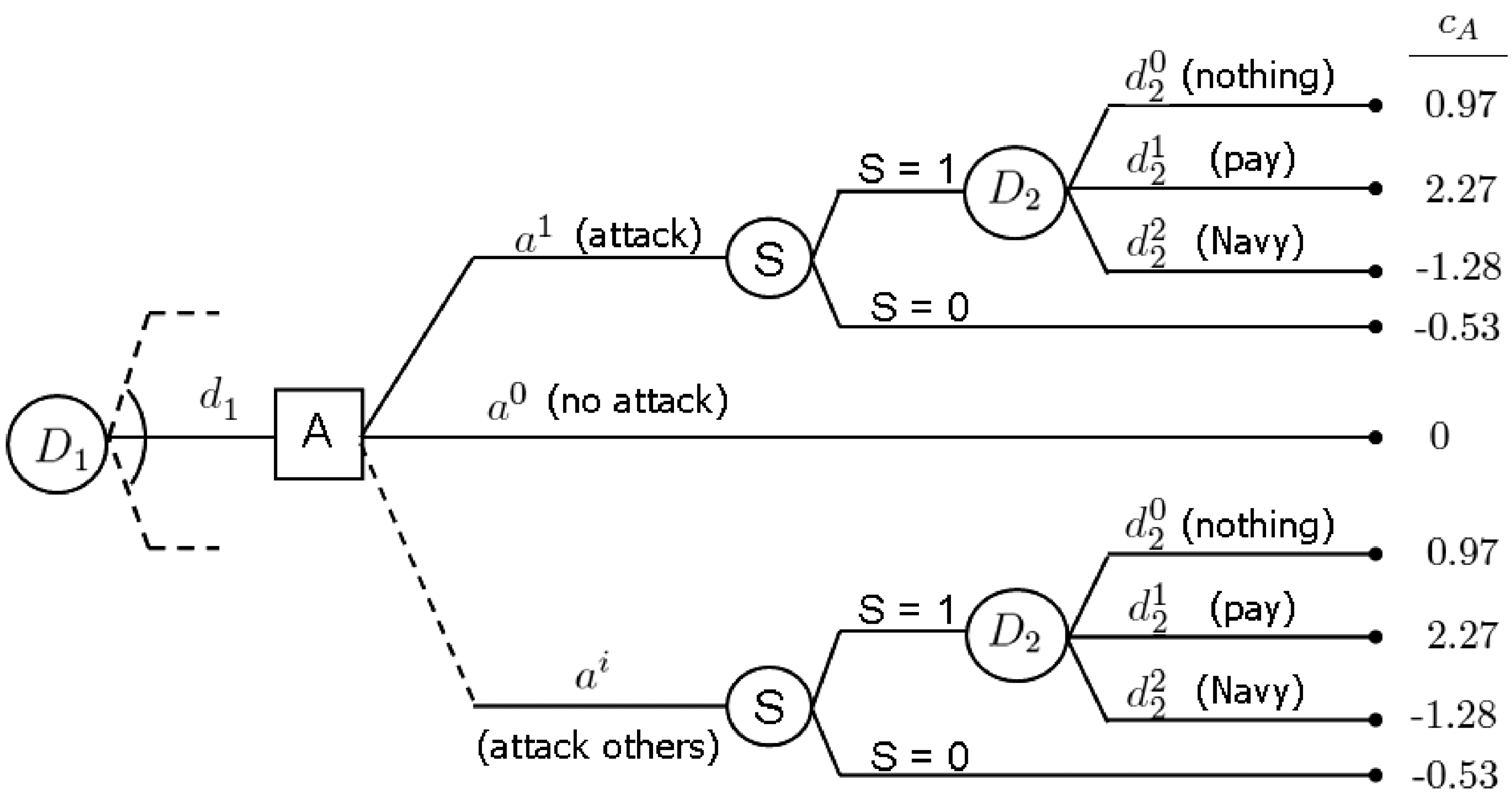

CASE STUDY: SOMALI PIRATES

- At decision node D_2, the Owner finds her maximum utility action, conditional on each $d_1 \in \mathcal{D}_1 \setminus \{d_1^3\}$, given by $d_2^*(d_1, a^1, S = 1) = \arg\max_{d_2 \in \mathcal{D}_2} u_D(c_D(d_1, S = 1, d_2))$

- Next, at chance node S she obtains her expected utility as

$$\begin{aligned}\Psi_D(d_1, a^1) &= p_D(S = 1|d_1, a^1)u_D(c_D(d_1, S = 1, d_2^*(d_1, a^1, S = 1))) + \\ &\quad + p_D(S = 0|d_1, a^1)u_D(c_D(d_1, S = 0))\end{aligned}$$

- At this point, she uses her assessments of the estimated probability of being attacked, conditional on her initial defense decision, or $\widehat{p}_D(A = a^1|d_1)$, to compute for each $d_1 \in \mathcal{D}_1 \setminus \{d_1^3\}$ her expected utility at chance node A:
  $\Psi_D(d_1) = \Psi_D(d_1, a^1)\widehat{p}_D(A = a^1|d_1) + u_D(c_D(d_1, S = 0))(1 - \widehat{p}_D(A = a^1|d_1))$

- Finally, she finds her maximum expected utility decision at decision node D_1 as $d_1^* = \arg\max_{d_1^i \in \mathcal{D}_1} \Psi_D(d_i^i)$, where $\Psi_D(d_1^3) = u_D(c_D(d_1^3))$ is obtained from the table with the Owner's costs (see next slide)

- The Defender's best strategy is to first choose d_1^* at node D_1, and, if the ship is hijacked, respond by choosing $d_2^*(d_1^*, a^1, S = 1)$ at node D_2

CASE STUDY: SOMALI PIRATES

e 4.2 The Owner's costs associated with different tree paths.

D_1	S	D_2	Ship loss	Action costs	Lives lost	c_D
do nothing	$S=1$	don't pay	1	0 + 0	0 + 4	15.16
do nothing	$S=1$	pay ransom	0	0 + 2.3M	0 + 0	2.30
do nothing	$S=1$	call Navy	0	0 + 0.2M	0 + 2	4.28
do nothing	$S=0$		0	0	0	0.00
hire guard	$S=1$	don't pay	1	0.05M + 0	1 + 4	17.25
hire guard	$S=1$	pay ransom	0	0.05M + 2.3M	1 + 0	4.39
hire guare	$S=1$	call Navy	0	0.05M + 0.2M	1 + 2	6.37
hire guard	$S=0$		0	0.05M	0	0.05
hire team	$S=1$	don't pay	1	0.15M + 0	2 + 4	19.39
hire team	$S=1$	pay ransom	0	0.15M + 2.3M	2 + 0	6.53
hire team	$S=1$	call Navy	0	0.15M + 0.2M	2 + 2	8.51
hire team	$S=0$		0	0.15M	0	0.15
d_1^3 (alternative route)			0	0.5M	0	0.50

CASE STUDY: SOMALI PIRATES

- This case study made many assumptions and a sensitivity analysis can show their impact
- First, recall that the Owner has a constant level of risk aversion, with utility function of the form $u_D(c_D) = -\exp(c \times c_D)$, for $c > 0$
- Then, if there are nine ships that may be targeted by Pirates, the Owner's maximum expected utility choice will be:
  - hire one armed guard for $c = 0.1$ and $c = 0.4$
  - hire two armed guards for $c = 1$ and $c = 2$
  - take the Cape of Good Hope route for $c = 5$
- In the first two cases, if hijacked, she should then pay the ransom

CASE STUDY: SOMALI PIRATES

- To also explore the sensitivity of the Owner's optimal decision to assumptions about (a) the probability that an attack is successful given the initial defense decision, (b) the number of nearby ships that could be targeted, and (c) the probability of an attempted attack given the initial defense decision, consider the next tables, showing the Owner's optimal initial and secondary decisions

- If hijacked, it is always best for the Owner to pay the ransom

- As the number of ships increases, risk diminishes and less aggressive initial choices are made

- Similarly, as the estimates of attack probabilities diminish, or as the probability of successful attack diminishes, the Owner should choose less expensive initial defenses

Table 4.4 A sensitivity analysis of the decision theory in the Somali pirates case study. Factors considered are different levels of risk aversion, different probabilities for successful attack, different numbers of nearby ships, and different probabilities of a successful attack given the initial decision.

$p_D(S=1\|a^1,d_1)$			$\hat{p}_D(A=a^1\|d_1)$					
d_1^1	d_1^2	n	d_1^0	d_1^1	d_1^2	c	d_1^*	$d_2^*(d_1^*)$
0.2	0.1	5	0.41010	0.18354	0.00382	$0.1-1$	d_1^2 (team)	d_2^1 (pay)
						$2-5$	d_1^3 (GH route)	
		9	0.27260	0.05780	0.00008	0.1	d_1^1 (man)	d_2^1 (pay)
						$0.4-1$	d_1^2 (team)	d_2^1 (pay)
						$2-5$	d_1^3 (GH route)	
		15	0.18322	0.01622	0.00000	$0.1-0.4$	d_1^1 (man)	d_2^1 (pay)
						$1-5$	d_1^2 (team)	d_2^1 (pay)
		20	0.14230	0.00628	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)
0.10	0.05	5	0.46564	0.12166	0.00328	0.1	d_1^1 (man)	d_2^1 (pay)
						$0.4-1$	d_1^2 (team)	d_2^1 (pay)
						$2-5$	d_1^3 (GH route)	
		9	0.30332	0.02560	0.00004	$0.1-0.4$	d_1^1 (man)	d_2^1 (pay)
						$1-2$	d_1^2 (team)	d_2^1 (pay)
						5	d_1^3 (GH route)	
		15	0.19386	0.00392	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)
		20	0.14836	0.00098	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)

Table 4.4 A sensitivity analysis of the decision theory in the Somali pirates case study. Factors considered are different levels of risk aversion, different probabilities for successful attack, different numbers of nearby ships, and different probabilities of a successful attack given the initial decision.

$p_D(S=1\|a^1,d_1)$			$\hat{p}_D(A=a^1\|d_1)$					
d_1^1	d_1^2	n	d_1^0	d_1^1	d_1^2	c	d_1^*	$d_2^*(d_1^*)$
0.10	0.05	5	0.46564	0.12166	0.00328	0.1	d_1^1 (man)	d_2^1 (pay)
						$0.4-1$	d_1^2 (team)	d_2^1 (pay)
						$2-5$	d_1^3 (GH route)	
		9	0.30332	0.02560	0.00004	$0.1-0.4$	d_1^1 (man)	d_2^1 (pay)
						$1-2$	d_1^2 (team)	d_2^1 (pay)
						5	d_1^3 (GH route)	
		15	0.19386	0.00392	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)
		20	0.14836	0.00098	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)
0.05	0.025	5	0.49010	0.09372	0.00374	0.1	d_1^1 (man)	d_2^1 (pay)
						$0.4-1$	d_1^2 (team)	d_2^1 (pay)
						$2-5$	d_1^3 (GH route)	
		9	0.31764	0.01596	0.00002	$0.1-0.4$	d_1^1 (man)	d_2^1 (pay)
						$1-2$	d_1^2 (team)	d_2^1 (pay)
						5	d_1^3 (GH route)	
		15	0.19842	0.00142	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)
		20	0.14778	0.00024	0.00000	$0.1-1$	d_1^1 (man)	d_2^1 (pay)
						$2-5$	d_1^2 (team)	d_2^1 (pay)