

Real-world Games

Boi Faltings

Laboratoire d'Intelligence Artificielle
boi.faltings@epfl.ch
<http://moodle.epfl.ch/>

Games in the real world

- Agents do not know each others' payoff matrices.
- Agents can cooperate to do better than a Nash equilibrium.
- Agents learn to select the best strategies, or...
- Agents negotiate to select their strategies.

Auction game

- auction for one single item. Agent A bids one of $\{1.5, 2.5, 3.5\}$, agent B bids one of $\{1, 2, 3\}$.
- the agent with the higher bid gets the item and pays its bid.
- let the value of the item be $\text{value}(A)=3$, $\text{value}(B)=2.5$

⇒ game:

		B		
		1	2	3
A	1.5	(1.5, 0)(A)	(0, 0.5)(B)	(0, -0.5)(B)
	2.5	(0.5, 0)(A)	(0.5, 0)(A)	(0, -0.5)(B)
	3.5	(-0.5, 0)(A)	(-0.5, 0)(A)	(-0.5, 0)(A)

NE: A plays 2.5, B plays 2

...but A and B do not know each other's values!

Games with uncertain utilities

Many games have uncertain utilities, for example trading or auctions:

- utility for each agent depends on its value for the item.
- this is *private* information.
- agents *type*: all information that only the agent knows.
- probability distribution of other agents' types is common knowledge.

How can agents play an equilibrium when utilities are not known?

Uncertain utilities

3 different ways of computing strategies with uncertain utility:

- Ex ante: assumes no knowledge of any agent's type.
(what is known *before* the game even starts)
- Ex interim: assumes knowledge of own type.
(what is known during the game)
- Ex post: assumes knowledge of all agents' types.
(what will be known in hindsight *after* the game)

Bayes-Nash equilibrium

Bayes-Nash equilibrium: Nash equilibrium in game with ex ante expected utilities.

Example: assume $p(1) = 0$, $p(2) = p(3) = 1/2$ for both A & B.
 $E[\text{value}(A)] = E[\text{value}(B)] = p(1) \cdot 1 + p(2) \cdot 2 + p(3) \cdot 3 = 2.5$.
⇒ expected game:

		B		
		1	2	3
A	1.5	(1,0)(A)	(0,0.5)(B)	(0,-0.5)(B)
	2.5	(0,0)(A)	(0,0)(A)	(0,-0.5)(B)
	3.5	(-1,0)(A)	(-1,0)(A)	(-1,0)(A)

(weakly) dominated actions: A=3.5, A=2.5, B=3

Bayes-Nash Equilibrium: A plays 1.5, B plays 2.

Ex-post Nash equilibrium ?

- Ex-post Nash equilibrium: strategies that gives the highest utilities no matter what the uncertain information will turn out to be.
- Does not necessarily exist: strategies may be different depending on other agents' types.
- Auction game: equilibrium is for A to bid $\min(\text{value}(A), \text{bid}(B) + 1)$
- But $\text{bid}(B)$ changes with $\text{value}(B)$, so A's best strategy can not be the same for all $\text{value}(B)$.

Ex-post Nash equilibrium

- Consider auction rule: winner pays *second highest* price.
- Claim: bidding true value is an ex-post Nash equilibrium.
- Assume $bid(B) = value(B)$. Then only two payoffs for A:
 - $bid(A) > value(B)$ have the same payoff $value(A) - value(B)$,
 - $bid(A) \leq value(B)$ have payoff zero.

Now consider the cases:

- $value(A) > value(B)$: payoff of $bid(A) = value(A)$ is > 0 : best response.
- $value(A) \leq value(B)$: payoff of $bid(A) > value(B)$ is < 0 ; $bid(A) = value(A)$ is a best response.
- Same reasoning for B .

General-sum games

- In general-sum games, agents should cooperate to obtain a higher payoff.
- Cooperation may not be a Nash equilibrium \Rightarrow players need to cooperate to achieve the best result.
- Joint plan and payoffs can be fixed by a contract that punishes deviation.
- Agents have to *negotiate* to agree on a joint strategy.

Prisoner's dilemma

- 2 suspects are arrested after a bank robbery and questioned (individually) by the police.
- Actions: choose between
 - cooperation (with the other suspect): deny all involvement in the crime.
 - defection: blame the other suspect for the crime.
- Knowledge:
A et B don't know the other's choice!
- Payoffs:
if A and B both cooperate, they are held by police for 1 year, and then can go off to enjoy their loot (utility 9). If both defect, they get 5 years in prison before (utility 5). If only one cooperates, he gets 10 years in prison (utility 0) while the other goes free (utility 10).

Business version

- 2 partners each put in 5 CHF in a joint effort
- Actions: choose between
 - cooperation: carry out the business together and each gain 9 CHF if successful.
 - defection: take the money and disappear.
- If both defect, they just get their money back.
- This is a very common business scenario.

Strategies

		B	
	C	C	D
A	D	(9,9) (0,10)	(10,0) (5,5)

Choice:

- cooperate: possible payoff = 0 or 9.
- defect: possible payoff = 5 or 10.

⇒ dominant strategies: both defect

Local and global optimality

- Dominant strategies: both players defect and get 5
- However, if both would agree to cooperate, the gain could be 9 for both of them.
- Not an equilibrium: either player can increase its gain from 9 to 10 by changing strategy.

⇒ requires a *contract* between players so that defection carries a punishment > 1 .

Mediated Equilibrium

Assume we have a *mediator*:

- agents can ask the mediator to play or play themselves.
- the mediator plays a known strategy as a function of the agents who asked it to play.
- mediator can be a vehicle to enforce a contract.

Example: Prisoner's dilemma

		B	
	C	C	D
C	(9,9)	(0,10)	
A	D	(10,0)	(5,5)

Dominant strategy equilibrium at (D,D)

Suppose a mediator plays:

- (C,C) if both players ask the mediator to play.
- D if only one of the players asks the mediator to play.

Prisoners' Dilemma with Mediator

		B		
		M	C	D
A	M	(9,9)	(10,0)	(5,5)
	C	(0,10)	(9,9)	(0,10)
	D	(5,5)	(10,0)	(5,5)

New dominant strategy equilibrium: (M,M)

Computers offer many possibilities to introduce mediators!

Correlated Equilibrium

Consider the "battle of the sexes":

		B
	O	S
O	(2,1)	(0,0)
A	S	(0,0)
		(1,2)

- 2 pure strategy Nash equilibria: (O,O) and (S,S): unfair!
- 1 mixed strategy Nash equilibrium: $([2/3, 1/3], [1/3, 2/3])$: fair, but expected payoff is only $4/3$.

Can we do better?

Correlated Equilibrium

Assume that there is a “trusted” coordinator that proposes to each agent i a choice of strategy s_i .

(the player does not have to follow the suggestion)

Original definition:

A correlated equilibrium *is a set of strategies $\{s_i\}$ such that for each agent i , choosing s_i as suggested by the coordinator is a best response to the strategies of the other agents (S_{-i}) .*

Example:

- fair coin flip $\Rightarrow (O, O)$ or (S, S)
- Equilibrium for player to stay with suggested strategy.
 \Rightarrow correlated equilibrium with expected payoffs $(1.5, 1.5)$.

More complex situation

		B
	O	O S
O	(2,1)	(0,0)
A S	(0,0)	(1,2)

- Let signal be (O,O) , (S,S) , (S,O) each with probability $1/3$.
- When A is assigned O, B will play O for sure. \Rightarrow best
- If A is assigned S, B plays O or S with equal probability.
 \Rightarrow better to play O and get $1/2 \cdot 2$ rather than $1/2 \cdot 1$!
 \Rightarrow not a correlated equilibrium!

Choosing the mapping signal \rightarrow strategy

Suppose both players observe a binary random variable $r \in \{0, 1\}$ (for example, a coin flip) and choose mapping to strategies:

		B			
		always O	$0 \rightarrow O, 1 \rightarrow S$	$0 \rightarrow S, 1 \rightarrow O$	always S
		(2,1)	(1,0.5)	(1,0.5)	(0,0)
A	always O	(1,0.5)	(1.5,1.5)	(0,0)	(0.5,1)
	$0 \rightarrow O$	(1.5,1.5)	(0,0)	(1.5,1.5)	(0.5,1)
	$1 \rightarrow S$	(0,0)	(0.5,1)	(0.5,1)	(1,2)

\Rightarrow two fair pure-strategy Nash equilibria with payoff (1.5, 1.5)!

Latent Coordinator

- Suppose correlation signal is *latent*, i.e. players know its distribution but cannot observe it.
- ⇒ Bayesian game: signal value is unknown.
- Agents choose action that is best response to opponents' *observed* play.
- ⇒ equilibrium can be found through learning: always play best response to strategies observed from others.
- Much easier and realistic to reach than Nash equilibria!

No-Regret

- Let s denote a joint strategy vector, i.e. $s \in S = \times_i S_i$
- A sequence of plays $\{s^0, s^1, \dots, s^T\}$ is said to be *no-regret* for i iff:

$$\sum_{t=0}^T u_i(s^t) \geq \max_{x \in S_i} \sum_{t=0}^T u_i(x, s_{-i}^t)$$

- At least as good as any fixed strategy in hindsight!

Coarse Correlated Equilibrium

- A coarse correlated equilibrium is a probability distribution p over the strategy vectors such that $\forall i$

$$\sum_s p(s) u_i(s) \geq \max_{x \in S_i} \sum_s p(s) u_i(x, s_{-i})$$

\Rightarrow for all agents i , induces a sequence of plays that are no-regret.

Coarse Correlated Equilibria (Examples)

Correlation device samples with equal probability from distributions:

- $(O,O), (S,S)$: both play O, S with probability $1/2$ each.
Expected payoff: $(1.5, 1.5)$
Better than best fixed strategy for A (O , average payoff = 1).
Better than best fixed strategy for B (S , average payoff = 1).
- $(O,O), (S,S), (O,S), (S,O)$: both play O, S with probability $1/2$.
Expected payoff: $(0.75, 0.75)$
 A is better off by always playing O , B by always playing S .
 \Rightarrow not a CCE!
- $(O,O), (S,S), (O,S)$:
 A plays O with prob. $2/3$, B with prob. $1/3$.
Expected payoff: $(2/3, 2/3)$ = the mixed Nash equilibrium.

$CCE \supset CE$

		B		
		1	2	3
A	1	(1,1)	(-1,-1)	(0,0)
	2	(-1,-1)	(1,1)	(0,0)
	3	(0,0)	(0,0)	(-1.1,-1.1)

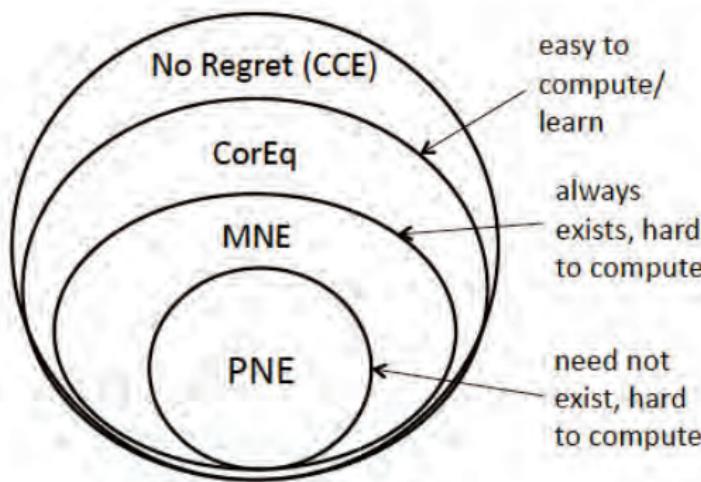
Consider playing (1,1), (2,2), (3,3) each with probability 1/3.

$$E[\text{payoff}] = 1/3 \cdot 1 + 1/3 \cdot 1 + 1/3 \cdot (-1.1) = 0.3$$

Best fixed actions (1 or 2): $E[\text{payoff}] = 0 \Rightarrow$ Coarse CE.

But not CE: when device suggests to play 3, agent is better off playing 1 or 2.

Hierarchy



The Price of Anarchy

- Explicit coordination: agents can coordinate on any strategy profile $\underline{s} \in \mathbf{S}$ with the highest joint reward $R(\underline{s})$.
- No coordination (anarchy): limited to equilibria $\underline{s}^i \in E$.
- Worst-case efficiency loss characterized by *Price of Anarchy*:

$$PoA = \frac{\max_{\underline{s} \in S} R(\underline{s})}{\min_{\underline{s} \in E} R(\underline{s})}$$

- Alternative for best-case: *Price of Stability*:

$$PoA = \frac{\max_{\underline{s} \in S} R(\underline{s})}{\max_{\underline{s} \in E} R(\underline{s})}$$

- Works for any kind of equilibria.

Bounding PoA

- can we bound PoA for a certain type of game?
- define: game with optimal strategy profile \underline{s}^* is (λ, μ) -smooth iff for every strategy profile \underline{s} :

$$\sum_{i \in \mathcal{A}} r_i(s_i^*, \underline{s}_{-i}) \geq \lambda R(\underline{s}^*) - \mu R(\underline{s})$$

⇒ PoA of a (λ, μ) -smooth game is at most $\lambda/(1 + \mu)$.

- many examples of smooth games: routing, facility location, simultaneous auctions, etc.

Improving beyond PoA

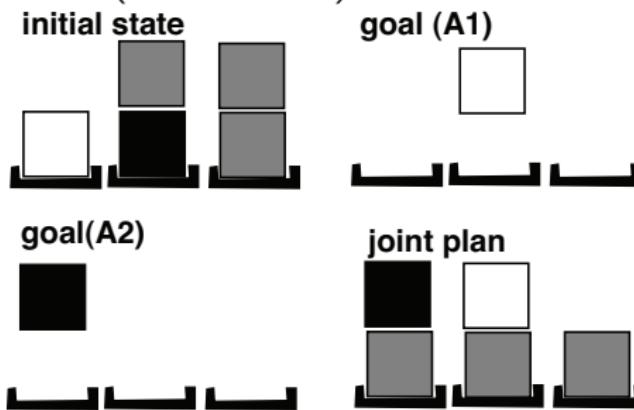
To implement a coordinated solution, we need:

- get agents to agree with prescribed strategy even when it is not a Nash equilibrium.
- find a solution that is fair to all agents.

⇒ negotiation to find an agreement.

Example: warehouse robots

Slotted blocks world (Rosenschein):



Each agent gets utility 5 when its goal is achieved.

Negotiation setting

One joint plan:

Agent A1 lifts white block, A2 moves three blocks around.

Cost: 2 operations (agent A1) vs. 6 operations (agent A2)

		B		
		N	A1	A2
A		N	(0,0)	(0,-2)
		A1	(-2,0)	(-2,-2)
		A2	(-6,0)	(-1,3)
				(-6,-6)

Equilibrium (conflict deal): (N,N)

Negotiation should reach (A1,A2) or (A2,A1) and side payment to compensate.

Example 2: Sharing wireless spectrum

- 2 agents A and B share a sequence of timeslots on a wireless channel to transmit sensor data.
- if they both transmit at the same time, most of transmission is lost (simultaneous defection).
- ⇒ use a time-division scheme so that A gets α and B $1 - \alpha$ of the slots.
- Strategies:
 - cooperate: agents transmit only in the assigned slots.
 - defect: agents transmit all the time.
- Defection is the dominant strategy (as in Prisoner's dilemma).
- Mediation requires agreement on α : negotiation.

Types of Negotiation

- Strategic negotiation: agents make and accept/reject offers in an unconstrained and self-interested manner.
- Axiomatic negotiation: agents agree on a set of axioms that the outcome should satisfy, then negotiate according to a protocol that guarantees the axioms.

Strategic negotiation

- Negotiation = sequence of rounds.
- Round: agent 1 makes an offer, agent 2 accepts or rejects.
- Next round: agent 2 makes offer, agent 1 accepts or rejects.
- Ends when an offer is accepted.

Alternating offers

- Consider scenario with 2 agents A and B.
- Protocol proceeds in n rounds:
 - 1 agent 1 makes a proposal P_1 for joint strategy $S(P_1)$ and payoffs $U_1(P_1), U_2(P_1)$.
 - 2 agent 2 accepts or rejects the proposal.where A and B take turns as agents 1 and 2.
- If negotiation fails, agents get conflict payoffs $U_1(C), U_2(C)$ = payoffs without coalition.
- Example: cutting a cake
 - Agent 1 proposes $\alpha \in [0..1]$, $U_1(\alpha) = \alpha, U_2(\alpha) = 1 - \alpha$
 - If no agreement, the cake is lost and both agents get 0.

Alternating offers with 1 round

- Assume selfish agents.
- Agent 2 accepts the offer P_1 iff $U_2(P_1) \geq U_2(C)$.
- Agent 1 should make an offer so that $U_1(P_1)$ is maximized and $U_1(P_1) \geq U_1(C), U_2(P_1) \geq U_2(C)$
- Best cake-cutting strategy for agent 1: propose $1 - \epsilon$.

Alternating offers with several rounds

- Let agent 1 be the one making the last offer.
- \Rightarrow in the last round, agent 1 can force any ϵ it wants!
- \Rightarrow agent 1 will not accept any offer of agent 2.
- All rounds before the last one are irrelevant!

Negotiation with time constraints

Suppose that the value of the cake decreases by factor δ_A for agent A and δ_B for agent B at each round.

- single round: agent 2 should accept anything.
- 2 rounds: agent 1 proposes $\alpha \leq 1 - \delta_2$, agent 2 accepts, because even if it got the whole cake in the next round, it would not get more utility than δ_2 which is already gets.
- many rounds: analyze as equilibrium.

Infinite duration with discount factors

- Agent A always offers x , agent B always offers y .
- Agent B should accept a offer that gives it at least $\delta_B y$:

$$(1 - x) \geq \delta_B y$$

- Symmetrically for agent A:

$$(1 - y) \geq \delta_A x$$

- Equilibrium: maximize shares \Rightarrow inequalities hold with equality:

$$x = \frac{1 - \delta_B}{1 - \delta_A \delta_B} \quad \text{and} \quad y = \frac{1 - \delta_A}{1 - \delta_A \delta_B}$$

- if $\delta_A = \delta_B = \delta$: $x = y = \frac{1}{1+\delta}$
- Agreement in the first step: maximizes joint return.

Problem with alternating offers

- In all cases, the agent who makes the first offer (agent A) gets a bigger share of the pie!
- Who decides who gets to make the first offer? Choice of protocol is not in equilibrium.
- More realistic:
 - both offers are made in parallel.
 - if they are not compatible, negotiation fails.
- What are the best strategies in such a game?

Framework for negotiation

- Agents have a set of *goals* $G = \{g_1, \dots, g_n\}$
- Agent i assigns each goal g a certain *worth* $w_i(g)$
- Agent i assigns each goal g a *standalone cost* $c_i^*(g)$
- *Deals* D_j are joint plans that achieve goals $G(D_j)$ at a certain cost $c_i(D_j)$ to agent i
- In the *conflict deal* D_c the agents do not cooperate and it has cost $c_i(D_c) = \sum_{g \in G(D_c)} c_i^*(g)$

Rational Action

Agents maximize their expected utility:

$$u_i(D_j) = \left[\sum_{g \in G(D_j)} w_i(g) \right] - c_i(D_j)$$

Agents do not have to cooperate: if negotiation does not succeed, they act independently and pursue the conflict deal.

Under what conditions is there a unique negotiation outcome?

Criteria for a negotiation outcome

Chosen deal \bar{D} should satisfy the criteria:

- feasible through a joint plan of action.
- pareto-optimal (non-dominated): there does not exist another deal D_k such that for all agents, $u_i(D_k) \geq u_i(\bar{D})$ and for at least one agent, $u_i(D_k) > u_i(\bar{D})$
- individually rational: for all agents, $u_i(\bar{D}) \geq u_i(D_c)$

Criteria for the solution (Nash)

3 more technical conditions for a unique solution:

- ① Feasibility.
- ② Pareto-Optimality.
- ③ Rationality.
- ④ Independence of sub-optimal alternatives:
If $\overline{D} \in T \subset S$, and \overline{D} is optimal within the results in S , then \overline{D} is optimal in T .
- ⑤ Independence of linear transformations:
If gains and losses are linearly transformed ($u' = \alpha u + \beta$), the new solution is the transformation of the old one.
- ⑥ Symmetry: If the game is symmetric for both players, then all agents get the same expected payoff.

Nash Bargaining Solution

- If there is a strategy \overline{D} that dominates D_c :
there is one single unique solution to the negotiation which
satisfies all 6 criteria.
- It is characterized by the condition:

$$(u_1(\overline{D}), \dots, u_n(\overline{D})) = \sup_D \prod_{i=1}^n (u_i(\overline{D}) - u_i(D_c))$$

where the maximization is carried out over all feasible deals.

- \Rightarrow provided that agents agree on the axioms, this is the
outcome of the negotiation!

Implementing the Nash Bargaining Solution

- A mediator collects all utilities and computes the Nash bargaining solution. But often no mediator (e.g. wireless spectrum)!
- Alternative without mediator:
 - 1 each agent A_i proposes a deal D_i .
 - 2 the plan that maximizes the product of agents' utility gains is chosen.

Each agent has an interest in proposing the best plan for everyone, since otherwise a suboptimal plan for itself might be chosen.

- Problem: every agent needs to know all others' utilities and strategies.

Reaching the Nash Solution by Alternating Offers

- Centralized mediation is very complex and requires detailed knowledge of all possible agent strategies.
- Q: Can we reach the Nash bargaining solution using agent-to-agent negotiation?
- A: yes, if agents follow certain rules.

Monotonic concession protocol (Zeuthen)

- Reach agreement through alternating offers.
- Offers from each agent must monotonically improve, i.e. agents progress by making *concessions*.
- Negotiation either ends when an offer is accepted, or fails when no agent has an interest to make further concessions.
- The agent that has the most to loose by negotiation failure has to make the next concession.

Risk indicators

- Suppose A_i rejects offer D_j and proposes D_i instead.
- This is rational only if:

$$u_i(D_j) - u_i(D_c) \leq p_i(u_i(D_i) - u_i(D_c))$$

- p_i = probability that negotiation will succeed in spite of rejecting D_j .
- Risk tolerance of A_i :

$$\begin{aligned} \text{risk}_i &= 1 - p_i^* \\ &= 1 - \frac{u_i(D_j) - u_i(D_c)}{u_i(D_i) - u_i(D_c)} = \frac{u_i(D_i) - u_i(D_j)}{u_i(D_i) - u_i(D_c)} \end{aligned}$$

$(p_i^* = \text{limit at equality})$

Monotonic concession protocol (Rosenschein)

Protocol:

- agents A_i, A_j both propose deals D_i, D_j .
- if one agrees to a proposal of the other, negotiation ends in agreement.
- otherwise, both calculate their risk tolerances $risk_i$ and $risk_j$; the agent with the smallest risk tolerance makes a concession.
- if none of the agents can rationally make a sacrifice, negotiation fails.

Limit case

- When $u_i(D_i) = u_i(D_c)$, $risk_i$ is undefined
- Agent A_i cannot make any further concessions without violating rationality! When should A_i make a concession?

$$risk_j = \frac{u_j(D_j) - u_j(D_i)}{u_j(D_j) - u_j(D_c)}$$

- If $risk_j > 1$, conflict deal offers better utility to A_j , so A_j should not make a concession and negotiation should end with conflict.
- If $risk_j < 1$, D_i is still more interesting to A_j so it should make a concession to approach it.

⇒ set $risk_i = 1$ to get the correct behavior

Properties of monotonic concessions

Smallest risk makes concession: eliminate deal D_i with largest p_i :

$$\frac{u_i(D_j) - u_i(D_c)}{u_i(D_i) - u_i(D_c)} > \frac{u_j(D_i) - u_j(D_c)}{u_j(D_j) - u_j(D_c)}$$
$$(u_i(D_j) - u_i(D_c))(u_j(D_j) - u_j(D_c)) > (u_i(D_i) - u_i(D_c))(u_j(D_i) - u_j(D_c))$$

- ⇒ maximizes product of utility gains
- ⇒ converges towards Nash solution!

Example: dividing wireless spectrum

- Agents goals: transmit packets of data.
- Utility for agent A: 3 per packet of data, for agent B: 9 per packet of data.
- Cost of transmission (for each packet): 1
- Conflict deal: both transmit their data all the time, success rate = 10% \Rightarrow payoff = (-0.7, -0.1).
- Goal of negotiation: decide $\alpha \in [0..1]$ so that A uses α of the slots and B uses $1 - \alpha$.

Nash Bargaining Solution

- Utilities: $U_A(\alpha) = \alpha(3 - 1)$, $U_B(\alpha) = (1 - \alpha)(9 - 1)$
- Nash solution: maximize $(2\alpha - (-0.7))(8(1 - \alpha) - (-0.1))$
 $\Rightarrow \alpha = 10.6/32 = 0.33$
- Note: A gets a smaller share of the channel.
 \Rightarrow incentive to lie and declare higher value
(A declares 9 $\Rightarrow \alpha = 0.5$).
- Makes sense only if claims can be verified.

Fairness of the Nash Bargaining Solution

- Utility gains over the conflict deal:

$$U_A(0.33) - U_A(D_c) = 0.66 - (-0.7) = 1.36$$

$$U_B(0.33) - U_B(D_c) = 5.28 - (-0.1) = 5.38$$

B gains about 4 times as much as A, since
B's utility per slot ($9 - 1 = 8$) is 4 times that of A ($3 - 1 = 2$).

- Due to the scale-invariance property!

Proposals...

- Initial proposals: $D_A : \alpha = 1, D_B : \alpha = 0$.

⇒ risks:

- $u_A(D_B) = 0, u_A(D_A) = 2 \Rightarrow risk_A = 2/2.7 = 0.74$
- $u_B(D_A) = 0, u_B(D_B) = 8 \Rightarrow risk_B = 8/8.1 = 0.99$

⇒ A has smaller tolerance and makes a concession!

- Next proposals: $D_A : \alpha = 0.5, D_B : \alpha = 0$

⇒ risks:

- $u_A(D_B) = 0, u_A(D_A) = 2 \Rightarrow risk_A = 1/1.7 = 0.69$
- $u_B(D_A) = 4, u_B(D_B) = 8 \Rightarrow risk_B = 4/8.1 = 0.49$

⇒ B has smaller tolerance and makes a concession!

- Next proposals: $D_A : \alpha = 0.5, D_B : \alpha = 0.25$

⇒ risks:

- $u_A(D_B) = 0.5, u_A(D_A) = 1 \Rightarrow risk_A = 0.5/1.7 = 0.29$
- $u_B(D_A) = 4, u_B(D_B) = 6 \Rightarrow risk_B = 2/6.1 = 0.32$

Generalization to > 2 agents

- Nash bargaining solution generalizes to n agents: maximize product of all agents' utility gains.
- Zeuthen protocol hard to extend.
- Use Nash formula to compute which proposal has lowest product of utility gains and ask that agent to make a concession.

Summary

- Uncertainty:
 - ex-ante/ex-interim: Bayes-Nash equilibria.
 - ex-post: only exists in certain cases.
- Correlated and coarse correlated equilibria.
- The best coordinated strategies are often not equilibria \Rightarrow require agreement by agents to act other than self-interested.
- Alternating offers protocol.
- Nash bargaining solution, monotonic concession protocol.
- Incentives for lying.