Game Theory

José M Vidal

Department of Computer Science and Engineering University of South Carolina

January 29, 2010

Abstract

Standard, extended, and characteristic form games. Chapters 2 and 3.

Outline

John von Neumann

- Born in Hungary. Came to US in 1930 to be professor at Princeton University.
- Participated in the Manhattan project. Coined the term MAD.
- Wrote "Theory of games and economic behavior" with Morgernstern.

John Von Neumann 1903–1957.

"..made contributions to quantum physics, functional analysis, set theory, economics, computer science, topology, numerical analysis, hydrodynamics (of explosions), statistics and many other mathematical fields as one of world history's outstanding mathematicians."

John F. Nash

- Born in the Appalachian mountains of West Virginia to an EE and a teacher.
- His PhD thesis at Princeton, in 1950, presented what we now call the Nash equilibrium, for which he won a Nobel prize in Economics in 1994.
- Diagnosed with paranoid schizophrenia in 1958 and worked to cure it until the 1990s. Feeling better now.
- Invented game of Hex.
- See book and movie "A Beautiful Mind".

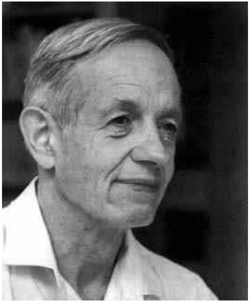

John F Nash, 1928

Outline

Outline

Payoff Matrix

		Alice	
		c	d
Bob	a	1,2	2,3
	b	4,5	6,7

- Payoff matrices represent the utility players can expect to receive given their choices.

Payoff Matrix

		Alice	
		c	d
Bob	a	1,2	2,3
	b	4,5	6,7

- Payoff matrices represent the utility players can expect to receive given their choices.
- Strategy s is set of actions players take.

Payoff Matrix

		Alice	
		c	d
Bob	a	1,2	2,3
	b	4,5	6,7

- Payoff matrices represent the utility players can expect to receive given their choices.
- Strategy s is set of actions players take.
  - They can be either pure or mixed.

Payoff Matrix

		Alice	
		c	d
Bob	a	1,2	2,3
	b	4,5	6,7

- Payoff matrices represent the utility players can expect to receive given their choices.
- Strategy s is set of actions players take.
  - They can be either pure or mixed.
- Players have common knowledge of the payoffs.

Payoff Matrix

		Alice	
		c	*d*
Bob	*a*	1,2	2,3
	b	4,5	6,7

- Payoff matrices represent the utility players can expect to receive given their choices.
- Strategy s is set of actions players take.
    - They can be either pure or mixed.
- Players have common knowledge of the payoffs.
- What should they do?

Assumptions and Requirements

- Players are rational (selfish). Participation is better than not.
- Strategy s is stable if no agent is motivated to diverge from it.
- A game is zero-sum if the sum of payoffs for every s is 0.

Outline

Solution Ideas

- Try to maximize your minimum utility: maxmin strategy.

Solution Ideas

- Try to maximize your minimum utility: maxmin strategy.
- The social welfare strategy is the one that maximizes the sum of everyone's payoffs.

Solution Ideas

- Try to maximize your minimum utility: maxmin strategy.
- The social welfare strategy is the one that maximizes the sum of everyone's payoffs.
- A strategy s is pareto optimal if there is no other strategy s' such that at least one agent is better off in s' and no agent is worse off in s' than in s.

Solution Ideas

- Try to maximize your minimum utility: maxmin strategy.
- The social welfare strategy is the one that maximizes the sum of everyone's payoffs.
- A strategy s is pareto optimal if there is no other strategy s' such that at least one agent is better off in s' and no agent is worse off in s' than in s.
- Strategy s is the dominant strategy for agent i if the agent is better off doing s no matter what the others do.

Solution Ideas

- Try to maximize your minimum utility: maxmin strategy.
- The social welfare strategy is the one that maximizes the sum of everyone's payoffs.
- A strategy s is pareto optimal if there is no other strategy s' such that at least one agent is better off in s' and no agent is worse off in s' than in s.
- Strategy s is the dominant strategy for agent i if the agent is better off doing s no matter what the others do.
    - In iterated dominance dominated strategies are eliminated in succession.

More Solution Ideas

- Strategy s is a Nash equilibrium if for all agents i, $s(i)$ is i's best strategy given that all the other players will play the strategies in s.
  - Nash showed that all game matrices have an equilibrium, but it might not be pure.

Maxmin Stratey

- Given by:

$$s_i^* = \max_{s_i} \min_{s_j} u_i(s_i, s_j). \tag{1}$$

Social Welfare Solution

- Agent i gets a utility $u_i(s_{-i}, s_i)$ when it takes action s_i and all others do s_{-i}.
- If we let $s = \{s_{-i}, s_i\}$ then we can say that the agent gets $u_i(s)$.
- The social welfare is

$$s^* = \arg\max_s \sum_i u_i(s)$$

Pareto Solution

- The pareto optimal is the set

$$\{s \mid \neg\exists_{s' \neq s}(\exists_i u_i(s') > u_i(s) \wedge \neg\exists_{j \in -i} u_j(s) > u_j(s'))\}$$

- Sometimes just called efficient.

Iterated Dominance

- A action a_i is dominant for agent i if

$$\forall_{a_{-i}} \forall_{b_i \neq a_i} u_i(a_{-i}, a_i) \geq u_i(a_{-i}, b_i)$$

- Apply repeatedly to all agents.

Iterated Dominance

- A action a_i is dominant for agent i if

$$\forall_{a_{-i}} \forall_{b_i \neq a_i} u_i(a_{-i}, a_i) \geq u_i(a_{-i}, b_i)$$

- Apply repeatedly to all agents.
- Might not reduce to one strategy.

Nash Equilibrium

- The set of strategies in Nash equilibrium is

$$\{s \mid \forall_i \forall_{a_i \neq s_i} u_i(s_{-i}, s_i) \geq u_i(s_{-i}, a_i)\}$$

Outline

Prisoner's Dilemma

Classic Prisoner's Dilemma

Two suspects A, B are arrested by the police. The police have insufficient evidence for a conviction, and having separated both prisoners, visit each of them and offer the same deal: if one testifies for the prosecution against the other and the other remains silent, the silent accomplice receives the full 10-year sentence and the betrayer goes free. If both stay silent, the police can only give both prisoners 6 months for a minor charge. If both betray each other, they receive a 2-year sentence each.

Prisoner's Dilemma

		A	
		Stays Silent	Betrays
B	Stays Silent	Both serve six months.	B serves 10 years; A goes free.
	Betrays	A serves 10 years; B goes free.	Both serve two years.

Canonical Prisoner's Dilemma

		A	
		Cooperate	*Defect*
B	*Cooperate*	3,3	0,5
	Defect	5,0	1,1

- Social Welfare =
- Pareto Optimal =
- Dominant =
- Nash =

Canonical Prisoner's Dilemma

		A	
		Cooperate	*Defect*
B	*Cooperate*	3,3	0,5
	Defect	5,0	1,1

- Social Welfare = (C,C)
- Pareto Optimal =
- Dominant =
- Nash =

Canonical Prisoner's Dilemma

		A	
		Cooperate	*Defect*
B	*Cooperate*	3,3	0,5
	Defect	5,0	1,1

- Social Welfare = (C,C)
- Pareto Optimal = (C,C) (D,C) (C,D)
- Dominant =
- Nash =

Canonical Prisoner's Dilemma

		A	
		Cooperate	*Defect*
B	*Cooperate*	3,3	0,5
	Defect	5,0	1,1

- Social Welfare = (C,C)
- Pareto Optimal = (C,C) (D,C) (C,D)
- Dominant = D for both players.
- Nash =

Canonical Prisoner's Dilemma

		A	
		Cooperate	*Defect*
B	*Cooperate*	3,3	0,5
	Defect	5,0	1,1

- Social Welfare = (C,C)
- Pareto Optimal = (C,C) (D,C) (C,D)
- Dominant = D for both players.
- Nash = (D, D)

Battle of the Sexes

Alice likes Ice hockey. Bob likes Football. They'd like to go out together. To which game does each one go?

Battle of the Sexes

		Alice	
		Ice Hockey	Football
Bob	Ice Hockey	4,7	0,0
	Football	3,3	7,4

- Social Welfare =
- Pareto Optimal =
- Dominant =
- Nash =

Battle of the Sexes

		Alice	
		Ice Hockey	Football
Bob	Ice Hockey	4,7	0,0
	Football	3,3	7,4

- Social Welfare = (I,I) (F,F)
- Pareto Optimal =
- Dominant =
- Nash =

Battle of the Sexes

		Alice	
		Ice Hockey	Football
Bob	Ice Hockey	4,7	0,0
	Football	3,3	7,4

- Social Welfare = (I,I) (F,F)
- Pareto Optimal = (I,I) (F,F)
- Dominant =
- Nash =

Battle of the Sexes

		Alice	
		Ice Hockey	Football
Bob	Ice Hockey	4,7	0,0
	Football	3,3	7,4

- Social Welfare = (I,I) (F,F)
- Pareto Optimal = (I,I) (F,F)
- Dominant = none.
- Nash =

Battle of the Sexes

		Alice	
		Ice Hockey	Football
Bob	Ice Hockey	4,7	0,0
	Football	3,3	7,4

- Social Welfare = (I,I) (F,F)
- Pareto Optimal = (I,I) (F,F)
- Dominant = none.
- Nash = (I,I) (F,F)

Chicken

Two maladjusted teenagers drive their cars towards each other at high speed. The one who swerves first is a chicken. If neither do, they both die.

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Social Welfare =
- Pareto Optimal =
- Dominant =
- Nash =

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Social Welfare = (C,S) (S,C)
- Pareto Optimal =
- Dominant =
- Nash =

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Social Welfare = (C,S) (S,C)
- Pareto Optimal = (C,S) (S,C)
- Dominant =
- Nash =

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Social Welfare = (C,S) (S,C)
- Pareto Optimal = (C,S) (S,C)
- Dominant = none.
- Nash =

Chicken

		Alice	
		Continue	Swerve
Bob	Continue	-1,-1	5,1
	Swerve	1,5	1,1

- Social Welfare = (C,S) (S,C)
- Pareto Optimal = (C,S) (S,C)
- Dominant = none.
- Nash = (C,S) (S,C)

Rational Pigs

There is one pig pen with a food dispenser at one end and the food comes out at the other end. It takes awhile to get from one side to the other. We put one big (strong) but slow pig, and a little, weak, and fast piglet. What happens?

Rational Pigs

		Pig	
		Nothing	Press Lever
Piglet	Nothing	0,0	5,1
	Press Lever	-1,6	1,5

- Social Welfare =
- Pareto Optimal =
- Dominant =
- Nash =

Rational Pigs

		Pig	
		Nothing	Press Lever
Piglet	Nothing	0,0	5,1
	Press Lever	-1,6	1,5

- Social Welfare = (N,P) (P,P)
- Pareto Optimal =
- Dominant =
- Nash =

Rational Pigs

		Pig	
		Nothing	Press Lever
Piglet	Nothing	0,0	5,1
	Press Lever	-1,6	1,5

- Social Welfare = (N,P) (P,P)
- Pareto Optimal = (N,P) (P,P) (P,N)
- Dominant =
- Nash =

Rational Pigs

		Pig	
		Nothing	Press Lever
Piglet	Nothing	0,0	5,1
	Press Lever	-1,6	1,5

- Social Welfare = (N,P) (P,P)
- Pareto Optimal = (N,P) (P,P) (P,N)
- Dominant = Piglet has N.
- Nash =

Rational Pigs

		Pig	
		Nothing	Press Lever
Piglet	Nothing	0,0	5,1
	Press Lever	-1,6	1,5

- Social Welfare = (N,P) (P,P)
- Pareto Optimal = (N,P) (P,P) (P,N)
- Dominant = Piglet has N.
- Nash = (N,P)

Outline

Iterated Games

- We let two players play the same game some number of times.

Iterated Games

- We let two players play the same game some number of times.
- Backward Induction: For any finite number of games defection is still the equilibrium strategy.
- However, practically we find that if there is a long time to go that people are more willing to cooperate.

Iterated Games

- We let two players play the same game some number of times.
- Backward Induction: For any finite number of games defection is still the equilibrium strategy.
- However, practically we find that if there is a long time to go that people are more willing to cooperate.
- A cooperative equilibrium can also be proven if instead of a fixed known number of interactions there is always a small probability that this will be the last interaction.

Folk Theorem

Theorem (Folk)

In a repeated game, any strategy where every agent gets a utility that is higher than his maxmin utility and is not Pareto-dominated by another is a feasible equilibrium strategy.

Folk Theorem

Theorem (Folk)

In a repeated game, any strategy where every agent gets a utility that is higher than his maxmin utility and is not Pareto-dominated by another is a feasible equilibrium strategy.

- Punish anyone who diverges by giving them their maxmin.
- It means: Much confusion.

Axelrod's Prisoner's Dilemma

- Robert Axelrod performed the now famous experiments on an iterated version of this problem.
- He sent out an email asking people to submit fortran programs that will play the PD against each other for 200 rounds. The winner was the one that accumulated more points.

Robert Axelrod

Iterated Prisoner's Dilemma Tournament

- ALL-D- always defect.
- RANDOM- pick randomly.
- TIT-FOR-TAT- cooperate in the first round, then do whatever the other player did last time.
- TESTER- defect first. If other player defects then play tit-for-tat. If he cooperated then cooperate for two rounds then defect.
- JOSS- play tit-for-tat but 10% of the time defect instead of cooperating.
- Which one won?

Iterated Prisoner's Dilemma Tournament

- ALL-D- always defect.
- RANDOM- pick randomly.
- TIT-FOR-TAT- cooperate in the first round, then do whatever the other player did last time.
- TESTER- defect first. If other player defects then play tit-for-tat. If he cooperated then cooperate for two rounds then defect.
- JOSS- play tit-for-tat but 10% of the time defect instead of cooperating.
- Which one won?
- Tit-for-tat won. It still made less than ALL-D when playing against it but, overall, it won more than any other strategy.
- Its was successful because it had the opportunity to play against other programs that were inclined to cooperate.

Axelrod's Lessons

- Do not be envious. You do not need to beat the other guy to do well yourself.
- Do not be the first to defect. This will usually have dire consequences in the long run.
- Reciprocate cooperation and defection. Not just one of them. You must reward and punish, with equal strengths.
- Do not be too clever. Trying to model what the other guy is doing leads you into infinite recursion since he might be modeling you modeling him modeling you.

Outline

Outline

Extended Form Game

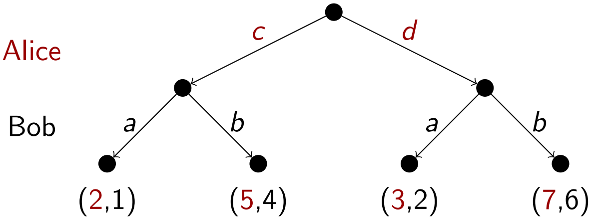

Extended Form Game

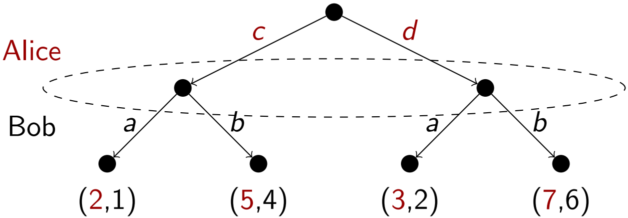

Outline

Subgame Perfect Equilibrium

The strategy s^* is a subgame perfect equilibrium if for all subgames, no agent i can get more utility than by playing s_i^* (assuming all others play s^*.

Multiagent MDPs

- Extended form games are nearly identical to multiagent MDPs.
- In practice, we use MMDPs.

Outline

Cooperative Games

- Mentioned in the original text, but not as popular (not mentioned in many introductory game theory textbooks).
- Model of the team formation problem.
  - Entrepreneurs trying to form small companies.
  - Companies cooperating to handle a large contract.
  - Professors colluding to write a grant proposal.

Outline

Formally, the General Characteristic Form Game

$A = \{1, \ldots, |A|\}$ the set of agents.

$\vec{u} = (u_1, \ldots, u_{|A|}) \in \Re^{|A|}$ is the outcome or solution.

$V(S) \subset \Re^{|S|}$ the rule maps every coalition $S \subset A$ to a utility possibility set.

Formally, the General Characteristic Form Game

$A = \{1, \ldots, |A|\}$ the set of agents.

$\vec{u} = (u_1, \ldots, u_{|A|}) \in \Re^{|A|}$ is the outcome or solution.

$V(S) \subset \Re^{|S|}$ the rule maps every coalition $S \subset A$ to a utility possibility set.

- For example, for the players $\{1, 2, 3\}$ we might have that $V(\{1, 2\}) = \{(5, 4), (3, 6)\}$.

Transferable Utility Game

- Assume that agents can freely trade utility.

Transferable Utility Game

- Assume that agents can freely trade utility.

Definition (Tranferable utility characteristic form game)

These games consist of a set of agents $A = \{1, \ldots, A\}$ and characteristic function $v(S) \rightarrow \Re$ defined for every $S \subseteq A$.

Transferable Utility Game

- Assume that agents can freely trade utility.

Definition (Tranferable utility characteristic form game)

These games consist of a set of agents $A = \{1, \ldots, A\}$ and characteristic function $v(S) \rightarrow \Re$ defined for every $S \subseteq A$.

- v is also called the value function.

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2 + 2 + 4 = 8$

(1)(23)
$2 + 8 = 10$

(2)(13)
$2 + 7 = 9$

(3)(12)
$4 + 5 = 9$

(123)
9

Outline

Feasibility

Definition (Feasible)

An outcome $\vec{u}$ is feasible if there exists a set of coalitions $T = S_1, \ldots, S_k$ where $\bigcup_{S \in T} S = A$ such that $\sum_{S \in T} v(S) \geq \sum_{i \in A} \vec{u}_i$.

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2 + 2 + 4 = 8$

(1)(23)
$2 + 8 = 10$

(2)(13)
$2 + 7 = 9$

(3)(12)
$4 + 5 = 9$

(123)
9

$\mathbf{u} = \{5, 5, 5\}$, is that feasible?

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2+2+4=8$

(1)(23)
$2+8=10$

(2)(13)
$2+7=9$

(3)(12)
$4+5=9$

(123)
9

$\mathbf{u} = \{5, 5, 5\}$, is that feasible? No

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2+2+4=8$

(1)(23)
$2+8=10$

(2)(13)
$2+7=9$

(3)(12)
$4+5=9$

(123)
9

$\mathbf{u} = \{2, 4, 3\}$, is that feasible?

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2 + 2 + 4 = 8$

(1)(23)
$2 + 8 = 10$

(2)(13)
$2 + 7 = 9$

(3)(12)
$4 + 5 = 9$

(123)
9

$\mathbf{u} = \{2, 4, 3\}$, is that feasible? Yes

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2 + 2 + 4 = 8$

(1)(23)
$2 + 8 = 10$

(2)(13)
$2 + 7 = 9$

(3)(12)
$4 + 5 = 9$

(123)
9

$\mathbf{u} = \{2, 2, 2\}$, is that feasible?

Example

S	$v(S)$
(1)	2
(2)	2
(3)	4
(12)	5
(13)	7
(23)	8
(123)	9

(1)(2)(3)
$2+2+4=8$

(1)(23)
$2+8=10$

(2)(13)
$2+7=9$

(3)(12)
$4+5=9$

(123)
9

$\mathbf{u} = \{2,2,2\}$, is that feasible? Yes,
but it is not stable.

The Core

Definition (Core)

An outcome $\vec{u}$ is in the *core* if

1. $$\forall_{S \subset A} : \sum_{i \in S} \vec{u}_i \geq v(S)$$
2. it is feasible.

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

$\vec{u} = \{2, 2, 2\}$ in core?

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

$\vec{u} = \{2, 2, 2\}$ in core? yes

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

$\vec{u} = \{2, 1, 2\}$ in core?

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1+2+2=5$

(1)(23)
$1+4=5$

(2)(13)
$2+3=5$

(3)(12)
$2+4=6$

(123)
6

$\vec{u} = \{2, 1, 2\}$ in core? no

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

$\vec{u} = \{1, 2, 2\}$ in core?

Example

S	$v(S)$
(1)	1
(2)	2
(3)	2
(12)	4
(13)	3
(23)	4
(123)	6

(1)(2)(3)
$1 + 2 + 2 = 5$

(1)(23)
$1 + 4 = 5$

(2)(13)
$2 + 3 = 5$

(3)(12)
$2 + 4 = 6$

(123)
6

$\vec{u} = \{1, 2, 2\}$ in core? no

Empty Cores Abound

S	$v(S)$
(1)	0
(2)	0
(3)	0
(12)	10
(13)	10
(23)	10
(123)	10

Good Definition, but

- In general, finding a solution in the core is not easy.

Lloyd Shapley

- How do we find an appropriate outcome?
- How do we fairly distribute the outcomes' value?
- What is fair?

Lloyd Shapley

- How do we find an appropriate outcome?
- How do we fairly distribute the outcomes' value?
- What is fair?

The Shapley value gives us one specific set of payments for coalition members, which are deemed fair.

Example

S	$v(S)$
()	0
(1)	1
(2)	3
(12)	6

- If they form (12), how much should each get paid?

Definition (Shapley Value)

Let $B(\pi, i)$ be the set of agents in the agent ordering π which appear before agent i. The *Shapley value* for agent i given A agents is given by

$$\phi(i, A) = \frac{1}{A!} \sum_{\pi \in \Pi_A} v(B(\pi, i) \cup i) - v(B(\pi, i)),$$

where Π_A is the set of all possible orderings of the set A. Another way to express the same formula is

$$\phi(i, A) = \sum_{S \subseteq A} \frac{(|A| - |S|)!\,(|S| - 1)!}{|A|!}[v(S) - v(S - \{i\})].$$

Example

S	$v(S)$
()	0
(1)	1
(2)	3
(12)	6

- If they form (12), how much should each get paid?

$$
\begin{aligned}
\phi(1,\{1,2\}) &= \frac{1}{2}\cdot(v(1)-v()+v(21)-v(2)) \\
&= \frac{1}{2}\cdot(1-0+6-3)=2 \\
\phi(2,\{1,2\}) &= \frac{1}{2}\cdot(v(12)-v(1)+v(2)-v()) \\
&= \frac{1}{2}\cdot(6-1+3-0)=4
\end{aligned}
$$

Drawbacks

- Requires calculating $A!$ orderings.
- Requires knowing $v(\cdot)$ for all coalitions.
- We still need to find the coalition structure.

Nucleolus

- Relax the core definition so that it will always exist.
- Idea: Find the solutions that minimizes the agents' temptation to defect.

Excess

Definition (excess)

The *excess* of coalition S given outcome $\vec{u}$ is given by

$$e(S, \vec{u}) = v(S) - \vec{u}(S),$$

where

$$\vec{u}(S) = \sum_{i \in S} \vec{u}_i.$$

Excess

Definition (excess)

The *excess* of coalition S given outcome $\vec{u}$ is given by

$$e(S, \vec{u}) = v(S) - \vec{u}(S),$$

where

$$\vec{u}(S) = \sum_{i \in S} \vec{u}_i.$$

The more excess S has, given $\vec{u}$, the more tempting it is for the agents in S to defect $\vec{u}$ and form S.

Nucleolus

Definition (nucleolus)

The *nucleolus* is the set

$$\{\vec{u} \,|\, \theta(\vec{u}) \not\succ \theta(\vec{v}) \text{ for all } \vec{v}, \text{ given that } \vec{u} \text{ and } \vec{v} \text{ are feasible.}\}$$

where,

$$\theta(\vec{u}) = \langle e(S_1^{\vec{u}}, \vec{u}), e(S_2^{\vec{u}}, \vec{u}), \ldots, e(S_{2^{|A|}}^{\vec{u}}, \vec{u})\rangle,$$

where $e(S_i^{\vec{u}}, \vec{u}) \geq e(S_j^{\vec{u}}, \vec{u})$ for all $i < j$.

Nucleolus

Definition (nucleolus)

The *nucleolus* is the set

$$\{\vec{u} \mid \theta(\vec{u}) \not\succ \theta(\vec{v}) \text{ for all } \vec{v}, \text{ given that } \vec{u} \text{ and } \vec{v} \text{ are feasible.}\}$$

where,

$$\theta(\vec{u}) = \langle e(S_1^{\vec{u}}, \vec{u}), e(S_2^{\vec{u}}, \vec{u}), \ldots, e(S_{2^{|A|}}^{\vec{u}}, \vec{u})\rangle,$$

where $e(S_i^{\vec{u}}, \vec{u}) \geq e(S_j^{\vec{u}}, \vec{u})$ for all $i < j$.
$\succ$ is a lexicographical ordering over all subsets S given some $\vec{u}$.
$\theta(\vec{u}) \succ \theta(\vec{v})$ is true when there is some number $q \in 1 \ldots 2^{|A|}$ such for all $p < q$ we have that $e(S_p^{\vec{u}}, \vec{u}) = e(S_p^{\vec{v}}, \vec{v})$ and $e(S_q^{\vec{u}}, \vec{u}) > e(S_q^{\vec{v}}, \vec{v})$ where the S_i have been sorted as per θ.

Lexicographic Example

For example, if we had the lists

$$\{(2,2,2),(2,1,0),(3,2,2),(2,1,1)\}$$

they would be ordered as

$$\{(3,2,2),(2,2,2),(2,1,1),(2,1,0)\}$$

.

Nucleolus

- Always exists.
- Captures idea of minimizing temptation, somewhat.
- Really, minimizes the greatest temptation.

Equal Excess

- Iterative algorithm for adjusting payments agents expect they will receive (adjust expectations).

Equal Excess

- Iterative algorithm for adjusting payments agents expect they will receive (adjust expectations).
- Let $E^t(i, S)$ be agent i's expected payoff for each coalition S which includes him.

Equal Excess

- Iterative algorithm for adjusting payments agents expect they will receive (adjust expectations).
- Let $E^t(i, S)$ be agent i's expected payoff for each coalition S which includes him.
- Let

$$A^t(i, S) = \max_{T \neq S} E^t(i, T)$$

  be agent i's expected payment from not choosing S and instead choosing the best alternative coalition.

Equal Excess

- Iterative algorithm for adjusting payments agents expect they will receive (adjust expectations).
- Let $E^t(i, S)$ be agent i's expected payoff for each coalition S which includes him.
- Let
$$A^t(i, S) = \max_{T \neq S} E^t(i, T)$$
be agent i's expected payment from not choosing S and instead choosing the best alternative coalition.
- Then, at each time step we update the players' expected payments using
$$E^{t+1}(i, S) = A^t(i, S) + \frac{v(S) - \sum_{j \in S} A^t(j, S)}{|S|}.$$

Outline

Centralized Algorithm: Search

(1)(2)(3)(4)

(12)(3)(4) (13)(2)(4) (14)(2)(3) (23)(1)(4) (24)(1)(3) (34)(1)(2)

(1)(234) (2)(134) (3)(124) (4)(123) (12)(34) (14)(23) (13)(24)

(1234)

Centralized Algorithm: Search

(1)(2)(3)(4)

(12)(3)(4) (13)(2)(4) (14)(2)(3) (23)(1)(4) (24)(1)(3) (34)(1)(2)

(1)(234) (2)(134) (3)(124) (4)(123) (12)(34) (14)(23) (13)(24)

(1234)

All possible coalitions

Search Order Bounds

Level	Bound
A	$A/2$
$A-1$	$A/2$
$A-2$	$A/3$
$A-3$	$A/3$
$A-4$	$A/4$
$A-5$	$A/4$
$\vdots$	$\vdots$
2	A
1	none

Distributed Search

FIND-COALITION(i)

$L_i \leftarrow$ set of all coalitions that include i.
$S_i^* \leftarrow \arg\max_{S \in L_i} v_i(S)$
$w_i^* \leftarrow v_i(S_i^*)$
Broadcast (w_i^*, S_i^*) and wait for all other broadcasts.
Put into W^*, S^* sets.
$w_{\max} = \max W^*$ and S_{max} is the corresponding coalition.
if $i \in S_{\max}$
    then join $S_{\max}$
Delete $S_{\max}$ from L_i.
Delete all $S \in L_i$ which include agents from $S_{\max}$.
if L_i is not empty
    then goto 2
return

Reduction to COP

- It can be reduced to a COP.

Outline

Coalition Formation

1. Agents generate values for the $v(\cdot)$ function.
2. Agents solve the characteristic form game by finding a suitable set of coalitions.
3. Agents distribute the payments from these coalitions to themselves in a suitable manner.