Strategic Games

Stéphane Airiau

Institute for Logic, Language & Computations (ILLC)
University of Amsterdam

12th European Agent Systems Summer School (EASSS 2010)
Ecole Nationale Supérieure des Mines de Saint-Etienne
Saint Etienne, France
August 23rd 2010.

- Agents have goals, they want to bring about some states of the world, they can take actions in their environment.
- In a multiagent system, agents interact, the actions of one may affect many other agents.
- How can we formally model such interactions?

Game theory is one way.

Outline

- **Today:** non-cooperative games
  - A central topic in Game theory: Strategic Games and Nash equilibrium.
  - Additional topics to provide a broader view of the field.
- **Tomorrow:** cooperative games

Prisoner's dilemma

Two partners in crime, Row (**R**) and Column (**C**), are arrested by the police and are being interrogated in separate rooms. From Row's point of view, four different outcomes can occur:

- only R confesses $\leadsto$ R gets 1 year.
- only C confesses $\leadsto$ R spends 4 years in jail
- both confess $\leadsto$ Both spend 3 years in prison.
- neither one confesses $\leadsto$ both get 2 years in prison

The utility of an agent is (5 - number of years in prison).

	Column confesses	Column does not
Row confesses	2,2	4,1
Row does not	1,4	3,3

We can abstract this game and provide a generic game representation as follows:

Definition (Normal form game)

A **normal form game (NFG)** is $(N, (S_i)_{i \in N}, (u)_{i \in N})$ where

- N is the set of n players
- S_i is the set of strategies available to agent i.
- $u_i : S_1 \times \cdots \times S_n \to \mathbb{R}^n$ is the **payoff function** of agent i. It maps a **strategy profile** to a **utility**.

Terminology:

- an element $s = \langle s_1, \ldots, s_n \rangle$ of $S_1 \times \cdots \times S_n$ is called a **strategy profile** or a **joint-strategy**.
- Let $s \in S_1 \times \cdots \times S_n$ and $s'_i \in S_i$. We write (s'_i, s_{-i}) the joint-strategy which is the same as s except for agent i which plays strategy s'_i, i.e., $(s'_i, s_{-i}) = \langle s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n \rangle$

What would you do?

- $N = \{Row, Column\}$
- $S_{Row} = S_{Column} = \{cooperate, defect\}$
- u_{Row} and u_{Column} are defined by the following bi-matrix.

Row \ *Column*	defect	cooperate
defect	2,2	4,1
cooperate	1,4	3,3

1. Wait to know the other action?
2. Not confess?
3. Confess?
4. Toss a coin?

Can you use some general principles to explain your choice?

Definition (strong dominance)

A strategy $x \in S_i$ for player i **(strongly) dominates** another strategy $y \in S_i$ if independently of the strategy played by the opponents, agent i (strictly) prefers x to y, i.e. $\forall s \in S_1 \times \cdots \times S_n$, $u_i(x, s_{-i}) > u_i(y, s_{-i})$

Prisoner's dilemma

	C confesses	C does not
R confesses	2,2	4,1
R does not	1,4	3,3

Both players have a dominant strategy: to confess! From Row's point of view:

- if C confesses: R is better off confessing as well.
- if C does not: R can exploit and confess.

Battle of the sexes

	L	R
T	2,2	4,3
B	3,4	1,1

- **Problem:** Where to go on a date: Soccer or Opera?
- **Requirements:**
  - have a date!
  - be at your favourite place!

Do players have a dominant strategy?

Definition (Best response)

A strategy s_i of a player i is a **best response** to a joint-strategy s_{-i} of its opponents iff

$$\forall s_i' \in S_i,\ u_i(s_i, s_{-i}) \geqslant u_i(s_i', s_{-i}).$$

Definition (Nash equilibrium)

A joint-strategy $s \in S_1 \times \cdots \times S_n$ is a **Nash equilibrium** if each s_i is a best response to s_{-i}, that is

$$(\forall i \in N)\ (\forall s_i' \in S_i)\ u_i(s_i, s_{-i}) \geqslant u_i(s_i', s_{-i})$$

Battle of the sexes possesses two Nash equilibria $\langle T, R\rangle$ and $\langle B, L\rangle$.

A **Nash equilibrium** is a joint-strategy in which no player could improve their payoff by unilaterally deviating from their assigned strategy.

Prisoner's dilemma

	C confesses	C does not
R confesses	2,2	4,1
R does not	1,4	3,3

Unique Nash equilibrium: both players confess!

- if R changes unilaterally, R loses!
- if C changes unilaterally, C loses!

Definition (Pareto optimal outcome)

A joint-strategy s is a **Pareto optimal outcome** if for no joint-strategy s'
$\forall i \in N\ u_i(s') \geqslant u_i(s)$ and $\exists i \in N u_i(s') > u_i(s)$

A joint-strategy is a Pareto optimal outcome when there is no outcome that is better for all players.

Prisoner's dilemma: Remaining silent is Pareto optimal.

discussion: It would be **rational** to confess! This seems counter-intuitive, as both players would be better off by keeping silent.

$\leadsto$ There is a conflict: the **stable** solution (i.e., the Nash equilibrium) is not **efficient**, as the outcome is not Pareto optimal.

Chicken game

In *Rebel Without a Cause*, James Dean's character's, Jim, is challenged to a "Chickie Run" with Buzz, racing stolen cars towards an abyss. The one who first jumps out of the car loses and is deemed a "chicken" (coward).

	Jim drives on	Jim turns
Buzz drives on	-10,-10	5,0
Buzz turns	0,5	1,1

Dominant Strategy?

Nash equilibrium ?

Nash equilibrium

- When there is no dominant strategy, an equilibrium is the next best thing.
- A game may not have a Nash equilibrium.
- If a game possesses a Nash equilibrium, it may not be unique.
- Any combinations of dominant strategies is a Nash equilibrium.
- A Nash equilibrium may not be Pareto optimal.
- Two Nash equilibria may not have the same payoffs

Definition (Mixed strategy)

A mixed strategy p_i of a player i is a probability distribution over its strategy space S_i.

Assume that there are three strategies: $S_i = \{1,2,3\}$. Player i may decide to play strategy 1 with a probability of $\frac{1}{3}$, strategy 2 with a probability of $\frac{1}{2}$ and strategy 3 with a probability of $\frac{1}{6}$. The mixed strategy is then denoted as $\left\langle \frac{1}{3}, \frac{1}{2}, \frac{1}{6} \right\rangle$.

Given a mixed strategy profile $p = \langle p_1, \ldots, p_n\rangle$, the expected utility for agent i is computed as follows:

$$E_i(p) = \sum_{s \in S_1 \times \cdots \times S_n} \left(\left(\prod_{j \in N} p_j(s_j) \right) \times u_i(s) \right)$$

Battle of the sexes

		y	$1-y$
		L	R
x	T	2,2	4,3
$1-x$	B	3,4	1,1

The expected utility for the Row player is:
$xy \cdot 2 + x(1-y) \cdot 4 + (1-x)y \cdot 3 + (1-x)(1-y) \cdot 1$
$= -4xy + 3x + 2y + 1$

Given a mixed strategy profile $p = \langle p_1, \dots, p_n \rangle$, we write (p'_i, p_{-i}) the mixed strategy profile which is the same as p except for player i which plays mixed strategy p'_i, i.e., $(p'_i, p_{-i}) = \langle p_1, \dots, p_{i-1}, p'_i, p_{i+1}, \dots, p_n \rangle$.

Definition (Mixed Nash equilibrium)

A **mixed Nash equilibrium** is a mixed strategy profile p such that $E_i(p) \geqslant E_i(p'_i, p_i)$ for every player i and every possible mixed strategy p'_i for i.

Battle of the sexes

	L	R
T	2,2	4,3
B	3,4	1,1

Let us consider that each player plays the mixed strategy $\langle \frac{3}{4}, \frac{1}{4} \rangle$.
None of the players have an incentive to deviate:

$E_{row}(T) = \frac{3}{4} \cdot 2 + \frac{1}{4} \cdot 4 = \frac{5}{2}$ $\quad E_{row}(B) = \frac{3}{4} \cdot 3 + \frac{1}{4} \cdot 1 = \frac{5}{2}$
(players are indifferent)

Theorem (J. Nash, 195))

Every finite strategic game has got at least one mixed Nash equilibrium.

note: The proofs are non-constructive and use Brouwer's or Kakutani's fixed point theorems.

J.F. Nash. Equilibrium points in n-person games. in *Proc. National Academy of Sciences of the United States of America*, 36:48-49, 1950.

Computing a Nash equilibrium

Complexity: In general, it is a hard problem. It is a PPAD-complete problem.

Daskalakis, Goldberg, Papadimitriou: **The complexity of computing a Nash equilibrium**, in *Proc. 38th Ann. ACM Symp. Theory of Computing (STOC)*, 2006

There are complexity results and algorithms for different classes of games. We will not treat then in this tutorial.

Y. Shoham & K. Leyton-Brown: **Multiagent Systems**, Cambridge University Press, 2009. (Chapter 4)
Nisan, Roughgarden, Tardos & Vazirani: **Algorithmic Game Theory**, Cambridge University Press, 2007. (chapters 2, 3)

Other types of solution concepts for NFGs

Safety strategy

With Nash equilibrium, we assumed that the opponents were **rational agents**. What if the opponents are potentially **malicious**, i.e., their goal could be to minimize the payoff of the player?

Definition (Maxmin)

For player i,
the **maxmin strategy** is $\arg\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$,
and its **maxmin value** or **safety level** is $\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$.

1) player i chooses a (possibly mixed) strategy.
2) the opponents $-i$ choose a (possible mixed) strategy that *minimize* i's payoff.
$\leadsto$ the maxmin strategy *maximizes* i's **worst case** payoff.

		y	$1-y$
		L	R
x	T	2,2	4,3
$1-x$	B	3,4	1,1

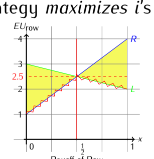

Payoff of Row
when Column plays pure strategy (T or R)
or any mixed strategy (yellow area)

Whatever Column does, Row can guarantee itself a payoff of 2.5 by playing the mixed strategy $\langle \frac{1}{2}, \frac{1}{2} \rangle$.

Punish

Definition (Minmax)

For player i in a 2-player game,
the **minmax strategy** is $\arg\min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$,
and its **minmax value** is $\min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$.

Player i's strategy against player $-i$ in a 2-player game is a strategy that minimizes $-i$'s best-case payoff

Proposition

For a player i,

$$\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \leqslant \min_{s_i \in S_i} \max_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})$$

Minimax theorem

Theorem
Minimax theorem (von Neumann, 1928)

In any finite two-player zero-sum game, for each player i, the maxmin strategy and minmax strategies are the same and are a Nash equilibrium of the game.

Minimax regret

Instead of assuming the opponents are rational (Nash equilibrium) or malicious (minimax), one can assume the **opponent is unpredictable** ⇝avoid **costly mistakes**/minimize their worst-case losses.

	L	R
T	100,100	0,0
B	0,0	1,1

(T,L) is preferred by both agents.
However, (B,R) is also a NE.
There is no dominance.
How to explain that (T,L) should be preferred?

One can build a **regret-recording** game where the payoff function r_i is defined by $r_i(s_i,s_{-i}) = u_i(s_i^\star,s_{-i}) - u_i(s_i,s_{-i})$, where $s_i^\star$ is i's best response to s_{-i}, i.e., $r_i(s_i,s_{-i})$ is i's **regret to have chosen s_i instead of $s_i^\star$**.

$r_i \backslash r_j$	L	R
T	0,0	1,100
B	100,1	0,0

We define $regret_i(s_i)$ as the maximal regret i can have from choosing s_i.
A **regret minimization strategy** is one that **minimizes the $regret_i$ function**.

Correlated equilibrium

Battle of the sexes

	L	R
T	2,2	4,3
B	3,4	1,1

How to avoid the bad outcomes in which the agents fail to coordinate?
idea: using a public random variable.

Example: the night before, the couple may condition their strategies based on weather (in the Netherlands, it is raining with a probability of 50%) as follows:
if it rains at 5pm, we go to opera, otherwise, we go to football.
⇝both players increase their expected utility
⇝maybe a fairer solution

Definition (Correlated equilibrium)
Given an n-agent game $G = (N,(S)_{i\in N},(u_i)_{i\in N})$, a correlated equilibrium is a tuple (v,π,σ), where

- v is a tuple of random variables $v = \langle v_1,\dots,v_n\rangle$ with respective domains $D = \langle D_1,\dots,D_n\rangle$,
- π is a joint-distribution over v,
- $\sigma = \langle \sigma_1,\dots,\sigma_n\rangle$ is a vector of mappings $\sigma_i : D_i \to S_i$,
- and for each agent i and every mapping $\sigma_i' : D_i \to S_i$ it is the case that $\sum_{d\in D}\pi(d)u_i(\sigma_1(d_1),\dots,\sigma_i(d_i),\dots,\sigma_n(d_n)) \geqslant \sum_{d\in D}\pi(d)u_i(\sigma_1(d_1),\dots,\sigma_i'(d_i),\dots,\sigma_n(d_n))$.

Theorem
For every Nash equilibrium, there exists a corresponding correlated equilibrium.

Proof
Let $s^\star$ be a Nash equilibrium. We define

- $D_i = S_i$: strategy space and the domains of the random variables are the same.
- $\pi(d) = \prod_{i\in N} s^\star(d_i)$
- $\sigma_i : D_i \to S_i$, $d_i \mapsto s_i$.

□

- Since a Nash equilibrium always exists, a correlated equilibrium always exists as well.
- However, a correlated equilibrium may not be a Nash equilibrium
- ⇝ correlated equilibrium is a generalization of Nash equilibrium.

- We have considered games where each player choose their action **simultaneously**, and we have studied the normal-form representation.
- They are many games which rely on turn-taking, e.g., chess, card games, etc. Game theory has something to say about these games as well.
- ⇝ We now introduce the **extended-form games (EFGs)**, in which a game is represented using a **tree** structure

Extended Form Games (EFGs)

Perfect-information game

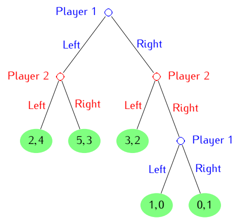

A game is described by a **game tree**.

- the leaf nodes contain the payoff to the agents.
- the non-leaf nodes are **choice nodes**, labeled with the agent that make the decision for the node.
- The game tree is **common knowledge** before the agents start to play.
- During the play, the agents know which actions have been chosen in the past: this is called the **perfect information** case.

A **strategy** is a complete plan of actions of a player: a strategy specifies an action for each of its choice node.

ex: Player 1 decides for two nodes and has four strategies: (Left, Left), (Left, Right), (Right, Left) and (Right, Right).

Perfect-information game

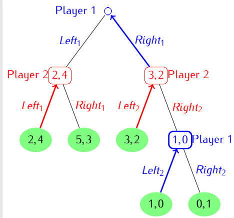

Backward induction: when an agent knows the payoff at each of a node's children, it can decide the best action of the player making the decision for this node.
If there are ties, then how they are broken affects what happens higher up in the tree
⇝Multiple equilibria...

From an EFG to a NFG

	L_1L_2	L_1R_1	R_1L_2	R_1R_2
L_1L_2	2,4	2,4	5,3	5,3
L_1R_2	2,4	2,4	5,3	5,3
R_1L_2	3,2	1,0	3,2	1,0
R_1R_2	3,2	0,1	3,2	0,1

- There can be an exponential number of pure strategies.
- Pure-strategy Nash equilibria of this game are (LL, LR), (LR, LR), (RL, LL), (RR, LL)
- But the only backward induction solution is (RL, LL)
- Nash equilibrium may be too weak for EFGs.

Definition (Subgame)

A **subgame** is any sub-tree of the game tree.

Definition (Subgame-perfect equilibrium)

A strategy profile s is a **subgame-perfect equilibrium** for an EFG G iff for any subgame g, the restriction of s to g is a Nash equilibrium of g.

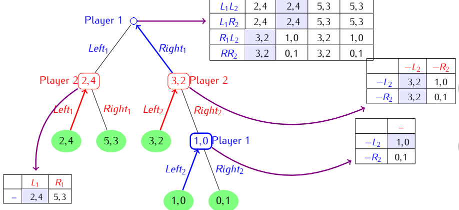

	L_1L_2	L_1R_2	R_1L_2	R_1R_2
L_1L_2	2,4	2,4	5,3	5,3
L_1R_2	2,4	2,4	5,3	5,3
R_1L_2	3,2	1,0	3,2	1,0
RR_2	3,2	0,1	3,2	0,1

	$-L_2$	$-R_2$
$-L_2$	3,2	1,0
$-R_2$	3,2	0,1

	$-$
$-L_2$	1,0
$-R_2$	0,1

	L_1	R_1
$-$	2,4	5,3

(L_1L_2, L_1R_2) ✘
(L_1R_2, L_1R_2) ✘
(R_1L_2, L_1L_2) ✔
(R_1R_2, L_1L_2) ✘

Other models of games

- **Congestion games:** a special game which always possess a **pure** strategy Nash equilibrium
- **Repeated games:** a NFG is played repeatedly (finitely/infinitely many times).
- **Stochastic games:** uncertainty about the next game to play
- **Bayesian games:** uncertainty about the current game

A **congestion game** is a tuple $(N, R, (S_i)_{i\in N}, (c_r)_{r\in R})$ where:

- $N = \{1, \ldots, n\}$ is the set of **players**
- $R = \{1, \ldots, m\}$ is the set of facilities or **resources**
- $S_i \subset M \setminus \emptyset$ denotes the set of **strategies** of player $i \in N$.
- $c_r(k)$ is the **cost** related to each user of resource $r \in M$ when exactly k players are using it.

Theorem

Every finite congestion game has a pure strategy Nash equilibrium.

R. W. Rosenthal. **A class of games possessing pure-strategy Nash equilibria**, in *International Journal of Game Theory*, 1973.

Theorem

Every congestion game is a potential game and every finite potential game is isomorphic to a congestion game

D. Monderer and L. S. Shapley **Potential Games**, in *Games and economic behavior*, 1996.

Repeated games

Prisoner's dilemma

	Defect	Cooperate
Defect	2,2	4,1
Cooperate	1,4	3,3

When players are **rational**, both players confess!
If they trusted each other, they could both not confess and obtain $\langle 3,3 \rangle$.
If the same players have to repeatedly play the game, then it could be rational not to confess.

- **One shot games**: there is no tomorrow.
  This is the type of games we have studied thus far.
- **Repeated games**: model a likelihood of playing the game again with the same opponent. The NFG (N, S, u) being repeated is called the **stage game**.
  - finitely repeated games ⇝ represent using a EFG and use backward induction to solve the game.
  - infinitely repeated games: the game tree would be infinite, use different techniques.

Infinitely repeated games

What is a strategy? In a repeated game, a **pure strategy** depends also on the **history** of play thus far.

ex: Tit-for-Tat strategy for the prisoner's dilemma:
Start by not confessing. Then, play the action played by the opponent during the previous iteration.

What is the players' objective?

- **Average criterion:** Average payoff received throughout the game by player i: $\lim_{t\to\infty} \frac{\sum_{t=1}^{k} u_i(s^t)}{k}$, where s^t is the joint-strategy played during iteration t.
- **Discounted-sum criterion:** Discounted sum of the payoff received throughout the game by player i: $\sum_{t=0}^{\infty} \gamma^t u_i(s^t)$, where γ is the discount factor (γ models how much the agent cares about the near term compared to long term).

Theorem (A Folk theorem)

Using the average criterion, any payoff vector v such that

- v is **feasible**, i.e., $\exists \lambda \in [0,1]^{\prod_{j\in N}|S_j|}$ s.t. $v_i = \sum_{s\in\prod_{j\in N}S_j} \lambda_s v_i(s)$
- v is **enforceable** $v_i \geqslant \max_{s_i\in S_i} \min_{s_{-i}\in S_{-i}} u_i(s_i, s_{-i})$

can be sustained by a Nash equilibrium.

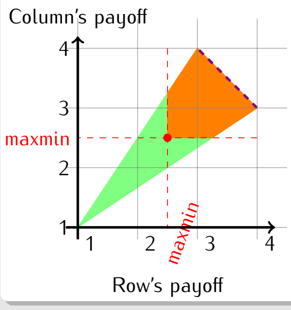

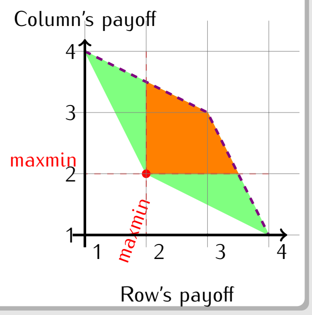

- In repeated games, the **same** *stage game* was played repeatedly.
- A **Stochastic game** is a set of NFGs. The agents **repeatedly** play games from this set. The next game is chosen with a probability which depends on the current game and the joint-action of the players.

Definition (Stochastic games)

A stochastic game is tuple $(N, (S_i)_{i\in N}, Q, P, (u_i)_{i\in N})$ where

- N is the set of players
- S_i is the strategy space of player i
- Q is a set of NFGs $q = (N, (S_i)_{i\in N}, (v_i^q)_{i\in N})$
- $P : Q \times \prod_{i\in N} S_i \times Q \to [0,1]$ is the **transition function**. $P(q, s, q')$ is the probability that game q' is played after game q when the joint-strategy s was played in game q.
- $u_i : Q \times \prod_{i\in N} S_i$ is the **payoff function**
  $u_i(q, s)$ is the payoff obtained by agent i when the joint-strategy s was played in game q.

In the definition, for ease of presentation, we assume that all the games have the same strategy space, which is not required.

- For stochastic games, the players know which game is currently played, i.e., they know the players of the game, the actions available to them, and their payoffs.
- In **Bayesian games**,
  - there is **uncertainty** about the game currently being played.
  - players have private information about the current game.
    The definition uses **information set**.

Definition (Bayesian game)

A **Bayesian game** is a tuple $(N, (S_i)_{i\in N}, G, P, (I_i)_{i\in N})$:

- N is the set of players.
- S_i is the set of strategies for agent i.
- G is a set of NFGs $g = (N, (S_i)_{i\in N}, (u_i^g)_{i\in N})$.
- P is a **common prior** over all games in G.
- I_i is the information set of agent i (a partition of G).
  A player knows the set which includes the current game, she does not know, however, which game it is in the set.
  ex: G is composed of six games, $I_2 = \{\{g_1, g_3, g_4\}, \{g_2, g_5\}\}$.
  Agent 2 knows the current game is in $\{g_1, g_3, g_4\}$, but she does not know whether the game is g_1, g_3, or g_4.

Evolutionary game theory

- models organisms in a large population (supposed infinite)
- **two** organisms are drawn randomly and play a 2-player game.
- the payoffs are linked to the fitness of the agents, and then, to their ability to reproduce.
- when an organism reproduces, a child adopt the same strategy as its parent.
- **Goal:** Are the strategies used by the organisms resilient to small mutant invasions? I.e, Is a strategy robust to evolutionary pressures?
  ⇝evolutionary stability.

J. W. Weibull, Evolutionary game theory, the MIT press, 1997

Summary and Concluding remarks

When does an agent play?

- Agents play **simultaneously** (Rock/Paper/Scissors) ⇝NFGs
- Agent play **sequentially** (chess, card games) ⇝EFGs

What is known?

- **Complete information** games: the structure of the game and the preference of the agents are common knowledge.
- **Incomplete information games**
  - does a player know the preference of its opponents? ⇝uncertainty, learning in games.
  - What kind of opponents? Rational? Malicious? ⇝Nash equilibrium, minmax, maxmin, regret.

What can be observed? Are the agents able to observe the actions of the opponents (**perfect/imperfect information**)

How does the game develop?

- Is it a one stage game?
- Are there multiple stages? (repeated games) Does the structure of the game change? ⇝Stochastic, Bayesian games
- Is the game played forever? ⇝Infinitely repeated games

Nobel Laureates

1972	Arrow	Social choice
1994	Nash, Selten and Harsanyi	Game theory
1996	Vickrey	Mechanism design
1998	Sen	Social choice
2005	Schelling and Aumann	Game theory
2007	Hurwicz, Maskin and Myerson	Mechanism design

- **Game theory:** mathematical study of interaction among independent, self-interested agents. (Two sessions at AAMAS-10)
  - non-cooperative games
  - cooperative games
  - games with sequential actions
  - evolutionary game theory
- **Mechanism design:** study of protocol design for strategic agents (one session at AAMAS-09)
- **Social choice:** study of preference aggregation / collective decision making. (One session at AAMAS-10)

Resources

- Martin J. Osborne and Ariel Rubinstein. **A course in Game Theory**, the MIT Press, 1994. (freely available online)
- Yoav Shoham and Kevin Leyton-Brown. **Multiagent Systems**, Cambridge University Press, 2009
- Michael Wooldridge. **An Introduction to Multiagent Systems**, Wiley, 2009
- Noam Nisan, Tim Roughgarden, Éva Tardos & Vijay V. Vazirani. **Algorithmic Game Theory**, Cambridge University Press, 2007.
- `gametheory.net`

Tomorrow

Cooperative games

When agents work together, the group of agents, as a whole, gets a payoff.

- What groups of agents to form?
- How to distribute the payoff to the individual agents?

Cooperative Games

Stéphane Airiau

Institute for Logic, Language & Computations (ILLC)
University of Amsterdam

12th European Agent Systems Summer School (EASSS 2010)
Ecole Nationale Supérieure des Mines de Saint-Etienne
Saint Etienne, France
August 23rd 2010.

Why study coalitional games?

Coalitional (or Cooperative) games are a branch of game theory in which **cooperation** or collaboration between agents can be modeled. Coalitional games can also be studied from a computational point of view (e.g., the problem of succint representation and reasoning).

A coalition may represent a set of:

- persons or group of persons (labor unions, towns)
- objectives of an economic project
- artificial agents

We have a population N of n agents.

Definition (Coalition)

A **coalition** $\mathcal{C}$ is a set of agents: $\mathcal{C} \in 2^N$.

Two main classes of games

1- Games with Transferable Utility (TU games)

- Two agents can **compare** their utility
- Two agents can **transfer** some utility

Definition (valuation or characteristic function)

A *valuation function* v associates a real number $v(S)$ to any subset S, i.e., $v : 2^N \to \mathbb{R}$

Definition (TU game)

A TU game is a pair (N, v) where N is a set of agents and where v is a valuation function.

2- Games with Non Transferable Utility (NTU games)

It is **not** always possible to compare the utility of two agents or to transfer utility (e.g., no price tags). Agents have preference over coalitions.

Informal example: a task allocation problem

- A set of tasks requiring different expertises needs to be performed, tasks may be decomposed.
- Agents do not have enough resource on their own to perform a task.
- Find complementary agents to perform the tasks
  - robots have the ability to move objects in a plant, but multiple robots are required to move a heavy box.
  - transportation domain: agents are trucks, trains, airplanes, ships... a task is a good to be transported.
- **Issues:**
  - What coalition to form?
  - How to reward each each member when a task is completed?

Some types of TU games

$\forall \mathcal{C}_1, \mathcal{C}_2 \subseteq N \mid \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$, $i \in N$, $i \notin \mathcal{C}_1$

- **additive (or inessential):** $v(\mathcal{C}_1 \cup \mathcal{C}_2) = v(\mathcal{C}_1) + v(\mathcal{C}_2)$ trivial from the game theoretic point of view
- **superadditive:** $v(\mathcal{C}_1 \cup \mathcal{C}_2) \geqslant v(\mathcal{C}_1) + v(\mathcal{C}_2)$ satisfied in many applications: it is better to form larger coalitions.
- **weakly superadditive:** $v(\mathcal{C}_1 \cup \{i\}) \geqslant v(\mathcal{C}_1) + v(\{i\})$
- **subadditive:** $v(\mathcal{C}_1 \cup \mathcal{C}_2) \leqslant v(\mathcal{C}_1) + v(\mathcal{C}_2)$
- **convex:** $\forall \mathcal{C} \subseteq \mathcal{T}$ and $i \notin \mathcal{T}$, $v(\mathcal{C} \cup \{i\}) - v(\mathcal{C}) \leqslant v(\mathcal{T} \cup \{i\}) - v(\mathcal{T})$.
  Convex game appears in some applications in game theory and have nice properties.
- **monotonic:** $\forall \mathcal{C} \subseteq \mathcal{T} \subseteq N$ $v(\mathcal{C}) \leqslant v(\mathcal{T})$.

The main problem

In the game (N, v) we want to form the **grand coalition**.

Each agent i will get a **personal payoff** x_i.

What are the interesting **properties** that x should satisfy?

How to **determine** the payoff vector x?

problem: a game (N, v) in which v is a worth of a coalition
solution: a payoff vector $x \in \mathbb{R}^n$

An example

$$N = \{1,2,3\}$$
$$v(\{1\}) = 0,\ v(\{2\}) = 0,\ v(\{3\}) = 0$$
$$v(\{1,2\}) = 90$$
$$v(\{1,3\}) = 80$$
$$v(\{2,3\}) = 70$$
$$v(\{1,2,3\}) = 105$$

What should we do?

- form $\{1,2,3\}$ and share equally $\langle 35,35,35 \rangle$?
- 3 can say to 1 "let's form $\{1,3\}$ and share $\langle 40,0,40 \rangle$".
- 2 can say to 1 "let's form $\{1,2\}$ and share $\langle 45,45,0 \rangle$".
- 3 can say to 2 "OK, let's form $\{2,3\}$ and share $\langle 0,46,24 \rangle$".
- 1 can say to 2 and 3, "fine! $\{1,2,3\}$ and $\langle 33,47,25 \rangle$
- ... is there a "good" solution?

Some properties

Let $x \in \mathbb{R}^n$ be a solution of the TU game (N,v)

Feasible solution: $\sum_{i \in N} x(i) \leqslant v(N)$.

Anonymity: a solution is independent of the names of the player.

Definition (marginal contribution)
The **marginal contribution** of agent i for a coalition $\mathcal{C} \subseteq N \setminus \{i\}$ is $mc_i(\mathcal{C}) = v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})$.

Let mc_i^{min} and mc_i^{max} denote the minimal and maximal marginal contribution.

x is **reasonable from above** if $\forall i \in N\ x^i < mc_i^{max}$

- ➪ mc_i^{max} is the strongest **threat** that an agent can use against a coalition.

x is **reasonable from below** if $\forall i \in N\ x^i > mc_i^{min}$

- ➪ mc_i^{min} is a minimum acceptable reward.

Some properties

Let x, y be two solutions of a TU-game (N,v).

Efficiency: $x(N) = v(N)$

- ➪ the payoff distribution is an allocation of the entire worth of the grand coalition to all agents.

Individual rationality: $\forall i \in N, x(i) \geqslant v(\{i\})$

- ➪ agent obtains at least its self-value as payoff.

Group rationality: $\forall \mathcal{C} \subseteq N, \sum_{i \in \mathcal{C}} x(i) = v(\mathcal{C})$

- ➪ if $\sum_{i \in \mathcal{C}} x(i) < v(\mathcal{C})$ some utility is lost.
- ➪ if $\sum_{i \in \mathcal{C}} x(i) > v(\mathcal{C})$ is not possible.

Pareto Optimal: $\sum_{i \in N} x(i) = v(N)$

- ➪ no agent can improve its payoff without lowering the payoff of another agent.

An **imputation** is a payoff distribution x that is efficient and individually rational.

The core

D Gillies, **Some theorems in n-person games.** *PhD thesis, Department of Mathematics, Princeton, N.J.*, 1953.

- A condition for a coalition to form:
  all agents prefer to be in it.
  i.e., none of the participants wishes she were in a different coalition or by herself ➪ **Stability**.
- Stability is a necessary but not sufficient condition, (e.g., there may be multiple stable coalitions).
- The **core** is a stability concepts where no agents prefer to deviate to form a different coalition.
- For simplicity, we will only consider the problem of the stability of the grand coalition:

➪ Is the grand coalition stable? $\Leftrightarrow$ Is the core non-empty?

The core relates to the stability of the grand coalition:
No group of agents has any incentive to change coalition.

Definition (*core* of a game (N,v))
Let (N,v) be a TU game, and assume we form the grand coalition N.
The **core** of (N,v) is the set:

$$Core(N,v) = \{x \in \mathbb{R}^n \mid x \text{ is a group rational imputation}\}$$

Equivalently,

$$Core(N,v) = \{x \in \mathbb{R}^n \mid x(N) \leqslant v(N) \wedge x(\mathcal{C}) \geqslant v(\mathcal{C})\ \forall \mathcal{C} \subseteq N\}$$

Weighted graph games

$$N = \{1,2\}$$
$$v(\{1\}) = 5,\ v(\{2\}) = 5$$
$$v(\{1,2\}) = 20$$

$core(N,v) = \{(x_1,x_2) \in \mathbb{R}^2 \mid x_1 \geqslant 5, x_2 \geqslant 5, x_1 + x_2 = 20\}$

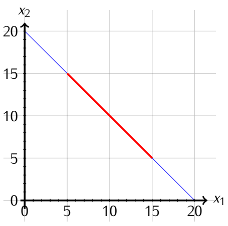

The core may not be fair: the core only considers stability.

Issues with the core

- The core may not always be non-empty.
- When the core is not empty, it may not be 'fair'.
- It may not be easy to compute.
- ➫ Are there classes of games that have a non-empty core?
- ➫ Is it possible to characterize the games with non-empty core?

Definition (Convex games)

A game (N,v) is **convex** iff
$\forall \mathcal{C} \subseteq \mathcal{T}$ and $i \notin \mathcal{T}$, $v(\mathcal{C} \cup \{i\}) - v(\mathcal{C}) \leqslant v(\mathcal{T} \cup \{i\}) - v(\mathcal{T})$.

TU-game is convex if the marginal contribution of each player increases with the size of the coalition he joins.

Theorem

A TU game (N,v) is convex iff for all coalition S and T
$v(S) + v(T) \leqslant v(S \cup T) + v(S \cap T)$

Theorem

A convex game has a non-empty core

Games with Coalition structures

Coalition Structure

Definition (Coalition Structure)

A **coalition structure (CS)** is a partition of the grand coalition into coalitions.
$\mathcal{S} = \{\mathcal{C}_1, \dots, \mathcal{C}_k\}$ where $\cup_{i \in \{1..k\}} \mathcal{C}_i = N$ and $i \neq j \Rightarrow \mathcal{C}_i \cap \mathcal{C}_j = \emptyset$.
We note $\mathscr{S}_N$ the set of all coalition structures over the set N.

ex: $\{\{1,3,4\}\{2,7\}\{5\}\{6,8\}\}$ is a coalition structure for $n = 8$ agents.

We start by defining a game with coalition structure, and see how we can define the core of such a game.

Game with Coalition Structure

Definition (TU game)

A TU game is a pair (N,v) where N is a set of agents and where v is a valuation function.

Definition (Game with Coalition Structures)

A **TU-game with coalition structure** $(N,v,\mathcal{S})$ consists of a TU game (N,v) and a CS $\mathcal{S} \in \mathscr{S}_N$.

- We assume that the players agreed upon the formation of $\mathcal{S}$ and only the payoff distribution choice is left open.
- The CS may model affinities among agents, may be due to external causes (e.g. affinities based on locations).
- The agents may refer to the value of coalitions with agents outside their coalition (i.e. opportunities they would have outside of their coalition).
- (N,v) and $(N,v,\{N\})$ represent the same game.

The set of **feasible** payoff vectors for $(N,v,\mathcal{S})$ is $X_{(N,v,\mathcal{S})} = \{x \in \mathbb{R}^n \mid \text{ for every } \mathcal{C} \in \mathcal{S} \ x(\mathcal{C}) \leqslant v(\mathcal{C})\}$.

Definition (Core of a game with CS)

The **core** $Core(N,v,\mathcal{S})$ of $(N,v,\mathcal{S})$ is defined by $\{x \in \mathbb{R}^n \mid (\forall \mathcal{C} \in \mathcal{S}, x(\mathcal{C}) \leqslant v(\mathcal{C})) \text{ and } (\forall \mathcal{C} \subseteq N, x(\mathcal{C}) \geqslant v(\mathcal{C}))\}$

We have $Core(N,v,\{N\}) = Core(N,v)$.

The next theorem is due to Aumann and Drèze.

R.J. Aumann and J.H. Drèze. **Cooperative games with coalition structures**, *International Journal of Game Theory*, 1974

Definition (Substitutes)

Let (N,v) be a game and $(i,j) \in N^2$. Agents i and j are **substitutes** iff $\forall \mathcal{C} \subseteq N \setminus \{i,j\}$, $v(\mathcal{C} \cup \{i\}) = v(\mathcal{C} \cup \{j\})$.

A nice property of the core related to fairness:

Theorem

Let $(N,v,\mathcal{S})$ be a game with coalition structure, let i and j be substitutes, and let $x \in Core(N,v,\mathcal{S})$. If i and j belong to different members of $\mathcal{S}$, then $x_i = x_j$.

The nucleolus

D. Schmeidler, **The nucleolus of a characteristic function game**. *SIAM Journal of applied mathematics*, 1969.

Excess of a coalition

Definition (Excess of a coalition)

Let (N,v) be a TU game, $\mathcal{C} \subseteq N$ be a coalition, and x be a payoff distribution over N. The **excess** $e(\mathcal{C},x)$ of coalition $\mathcal{C}$ at x is the quantity $e(\mathcal{C},x) = v(\mathcal{C}) - x(\mathcal{C})$.

An example: let $N = \{1,2,3\}$, $\mathcal{C} = \{1,2\}$, $v(\{1,2\}) = 8$, $x = \langle 3,2,5 \rangle$, $e(\mathcal{C},x) = v(\{1,2\}) - (x_1 + x_2) = 8 - (3+2) = 3$.

We can interpret a positive excess ($e(\mathcal{C},x) \geqslant 0$) as the amount of **dissatisfaction** or **complaint** of the members of $\mathcal{C}$ from the allocation x.

We can use the excess to define the core: $Core(N,v) = \{x \in \mathbb{R}^n \mid x \text{ is an imputation and } \forall \mathcal{C} \subseteq N, e(\mathcal{C},x) \leqslant 0\}$

This definition shows that no coalition has any complaint: each coalition's demand can be granted.

$$N = \{1,2,3\},\ v(\{i\}) = 0 \text{ for } i \in \{1,2,3\}$$
$$v(\{1,2\}) = 5,\ v(\{1,3\}) = 6,\ v(\{2,3\}) = 6$$
$$v(N) = 8$$

Let us consider two payoff vectors $x = \langle 3,3,2 \rangle$ and $y = \langle 2,3,3 \rangle$.
Let $e(x)$ denote the sequence of **excesses** of all coalitions at x.

$x = \langle 3,3,2 \rangle$

coalition $\mathcal{C}$	$e(\mathcal{C},x)$
{1}	-3
{2}	-3
{3}	-2
{1,2}	-1
{1,3}	1
{2,3}	1
{1,2,3}	0

$y = \langle 2,3,3 \rangle$

coalition $\mathcal{C}$	$e(\mathcal{C},y)$
{1}	-2
{2}	-3
{3}	-3
{1,2}	0
{1,3}	1
{2,3}	0
{1,2,3}	0

Which payoff should we prefer? x or y? Let us write the excess in the decreasing order (from the greatest excess to the smallest)

$\langle 1,1,0,-1,-2,-3,-3 \rangle$ $\langle 1,0,0,0,-2,-3,-3 \rangle$

Definition (lexicographic order of $\mathbb{R}^m$ $\geqslant_{lex}$)

Let $\geqslant_{lex}$ denote the **lexicographical** ordering of $\mathbb{R}^m$, i.e., $\forall (x,y) \in \mathbb{R}^m$, $x \geqslant_{lex} y$ iff

$$\begin{cases} \text{x=y or} \\ \exists t \text{ s. t. } 1 \leqslant t \leqslant m \text{ s. t. } \forall i \text{ s. t. } 1 \leqslant i \leqslant t \ x_i = y_i \text{ and } x_t > y_t \end{cases}$$

example: $\langle 1,1,0,-1,-2,-3,-3 \rangle \geqslant_{lex} \langle 1,0,0,0,-2,-3,-3 \rangle$

Let l be a sequence of m reals. We denote by $l^{\blacktriangleright}$ the **reordering** of l in **decreasing** order.

In the example, $e(x) = \langle -3,-3,-2,-1,1,1,0 \rangle$, then $e(x)^{\blacktriangleright} = \langle 1,1,0,-1,-2,-3,-3 \rangle$.

Hence, we can say that y is better than x by writing
$$e(x)^{\blacktriangleright} \geqslant_{lex} e(y)^{\blacktriangleright}.$$

Definition (Nucleolus)

Let (N,v) be a TU game.
Let $\mathcal{I}mp$ be the set of all imputations.
The **nucleolus** $Nu(N,v)$ is the set
$Nu(N,v) = \{x \in \mathcal{I}mp \mid \forall y \in \mathcal{I}mp \;\; e(y)^{\blacktriangleright} \succeq_{lex} e(x)^{\blacktriangleright}\}$

Theorem (Non-emptyness of the nucleolus)

Let (N,v) be a TU game, if $\mathcal{I}mp \neq \emptyset$,
then the nucleolus $Nu(N,v)$ is **non-empty**.

For a TU game (N,v) the nucleolus $Nu(N,v)$ is non-empty when $\mathcal{I}mp \neq \emptyset$, which is a great property as agents will always find an agreement. But there is more!

Theorem

The nucleolus has **at most one** element

In other words, there is **one** agreement which is stable according to the nucleolus.

Theorem

Let (N,v) be a superadditive game and $\mathcal{I}mp$ be its set of imputations. Then, $\mathcal{I}mp \neq \emptyset$.

Proof

Let (N,v) be a superadditive game.
Let x be a payoff distribution defined as follows:
$x_i = v(\{i\}) + \frac{1}{|N|}\left(v(N) - \sum_{j\in N} v(\{j\})\right)$.

- $v(N) - \sum_{j\in N} v(\{j\}) > 0$ since (N,v) is superadditive.
- It is clear x is individually rational ✔
- It is clear x is efficient ✔

Hence, $x \in \mathcal{I}mp$. □

Theorem

Let (N,v) be a TU game with a non-empty core. Then $Nu(N,v) \subseteq Core(N,v)$

The kernel.

M. Davis. and M. Maschler, **The kernel of a cooperative game**. *Naval Research Logistics Quarterly*, 1965.

Excess

Definition (Excess)

For a TU game (N,v), the excess of coalition $\mathcal{C}$ for a payoff distribution x is defined as $e(\mathcal{C},x) = v(\mathcal{C}) - x(\mathcal{C})$.

We saw that a positive excess can be interpreted as an amount of complaint for a coalition.
We can also interpret the excess as a potential to generate more utility.

Definition (Maximum surplus)

For a TU game (N,v), the **maximum surplus** $s_{k,l}(x)$ of **agent** k **over agent** l with respect to a payoff distribution x is the **maximum excess** from a coalition that **includes** k but does **exclude** l, i.e.,
$s_{k,l}(x) = \max_{\mathcal{C}\subseteq N \mid k\in\mathcal{C},\, l\notin\mathcal{C}} e(\mathcal{C},x)$.

Definition (Kernel)

Let $(N,v,\mathcal{S})$ be a TU game with coalition structure. The **kernel** is the set of imputations $x \in X_{(N,v,\mathcal{S})}$ such that for every coalition $\mathcal{C} \in CS$, if $(k,l) \in \mathcal{C}^2$, $k \neq l$, then we have either $s_{kl}(x) \geqslant s_{lk}(x)$ or $x_k = v(\{k\})$.

$s_{kl}(x) < s_{lk}(x)$ calls for a transfer of utility from k to l unless it is prevented by individual rationality, i.e., by the fact that $x_k = v(\{k\})$.

Properties

Theorem

Let $(N,v,\mathcal{S})$ a game with coalition structure, and let $\mathcal{I}mp \neq \emptyset$. Then we have $Nu(N,v,\mathcal{S}) \subseteq K(N,v,\mathcal{S})$

Theorem

Let $(N,v,\mathcal{S})$ a game with coalition structure, and let $\mathcal{I}mp \neq \emptyset$. The kernel $K(N,v,\mathcal{S})$ of the game is non-empty.

Proof

Since the Nucleolus is non-empty when $\mathcal{I}mp \neq \emptyset$, the proof is immediate using the theorem above. □

Computing a kernel-stable payoff distribution

- There is a transfer scheme converging to an element in the kernel.
- It may require an infinite number of small steps.
- We can consider the ϵ-kernel where the inequality are defined up to an arbitrary small constant ϵ.

R. E. Stearns. **Convergent transfer schemes for n-person games**. *Transactions of the American Mathematical Society*, 1968.

Computing a kernel-stable payoff distribution

Algorithm 1: Transfer scheme converging to a ϵ-Kernel-stable payoff distribution for the CS $\mathcal{S}$

```
compute-ε-Kernel-Stable(N, v, S, ε)
repeat
    for each coalition C ∈ S do
        for each member (i,j) ∈ C, i ≠ j do          // compute the maximum surplus
            // for two members of a coalition in S
            s_ij ← max_{R⊆N|(i∈R, j∉R)} v(R) − x(R)
    δ ← max_{(i,j)∈C², C∈S} s_ij − s_ji;
    (i*, j*) ← argmax_{(i,j)∈N²} (s_ij − s_ji);
    if (x_j* − v({j}) < δ/2) then                    // payment should be individually rational
        d ← x_j* − v({j*});
    else
        d ← δ/2;
    x_i* ← x_i* + d;
    x_j* ← x_j* − d;
until δ/v(S) ⩽ ε ;
```


- The complexity for one side-payment is $O(n \cdot 2^n)$.
- Upper bound for the number of iterations for converging to an element of the ϵ-kernel: $n \cdot log_2(\frac{\delta_0}{\epsilon \cdot v(S)})$, where δ_0 is the maximum surplus difference in the initial payoff distribution.
- To derive a polynomial algorithm, the number of coalitions must be bounded. For example, only consider coalitions which size is bounded in $[K_1, K_2]$. The complexity of the truncated algorithm is $O(n^2 \cdot n_{coalitions})$ where $n_{coalitions}$ is the number of coalitions with size in$[K_1, K_2]$, which is a polynomial of order K_2.

• M. Klusch and O. Shehory. **A polynomial kernel-oriented coalition algorithm for rational information agents**. In *Proceedings of the Second International Conference on Multi-Agent Systems*, 1996.
• O. Shehory and S. Kraus. **Feasible formation of coalitions among autonomous agents in non-superadditve environments**. *Computational Intelligence*, 1999.

The Shapley value

Lloyd S. Shapley. **A Value for *n*-person Games**. In *Contributions to the Theory of Games, volume II (Annals of Mathematical Studies)*, 1953.

Definition (marginal contribution)
The **marginal contribution** of agent i for a coalition $\mathcal{C} \subseteq N \setminus \{i\}$ is $mc_i(\mathcal{C}) = v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})$.

$\langle mc_1(\emptyset), mc_2(\{1\}), mc_3(\{1,2\}) \rangle$ is an efficient payoff distribution for any game $(\{1,2,3\}, v)$. This payoff distribution may model a dynamic process in which 1 starts a coalition, is joined by 2, and finally 3 joins the coalition $\{1,2\}$, and where the incoming agent gets its marginal contribution.

An agent's payoff depends on which agents are already in the coalition. This payoff may not be **fair**. To increase fairness,one could take the average marginal contribution over all possible joining orders.

Let σ represent a joining order of the grand coalition N, i.e., σ is a permutation of $\langle 1, \dots, n \rangle$.
We write $mc(\sigma) \in \mathbb{R}^n$ the payoff vector where agent i obtains $mc_i(\{\sigma(j) \mid j < i\})$. The vector mc is called a **marginal vector**.

Shapley value: version based on marginal contributions

Let (N, v) be a TU game. Let $\Pi(N)$ denote the set of all permutations of the sequence $\langle 1, \dots, n \rangle$.

$$Sh(N, v) = \frac{\sum_{\sigma \in \Pi(N)} mc(\sigma)}{n!}$$

the Shapley value is a **fair** payoff distribution based on marginal contributions of agents averaged over joining orders of the coalition.

An example

$N=\{1,2,3\}$, $v(\{1\})=0$, $v(\{2\})=0$, $v(\{3\})=0$,
$v(\{1,2\})=90$, $v(\{1,3\})=80$, $v(\{2,3\})=70$,
$v(\{1,2,3\})=120$.

	1	2	3
$1\leftarrow 2\leftarrow 3$	0	90	30
$1\leftarrow 3\leftarrow 2$	0	40	80
$2\leftarrow 1\leftarrow 3$	90	0	30
$2\leftarrow 3\leftarrow 1$	50	0	70
$3\leftarrow 1\leftarrow 2$	80	40	0
$3\leftarrow 2\leftarrow 1$	50	70	0
total	270	240	210
Shapley value	45	40	35

Let $y=\langle 50,40,30\rangle$

$\mathcal{C}$	$e(\mathcal{C},x)$	$e(\mathcal{C},y)$
{1}	-45	0
{2}	-40	0
{3}	-35	0
{1,2}	5	0
{1,3}	0	0
{2,3}	-5	0
{1,2,3}	120	0

This example shows that the Shapley value may not be in the core, and may not be the nucleolus.

- There are $|\mathcal{C}|!$ permutations in which all members of $\mathcal{C}$ precede i.
- There are $|N\setminus(\mathcal{C}\cup\{i\})|!$ permutations in which the remaining members succede i, i.e. $(|N|-|\mathcal{C}|-1)!$.

The Shapley value $Sh_i(N,v)$ of the TU game (N,v) for player i can also be written

$$Sh_i(N,v)=\sum_{\mathcal{C}\subseteq N\setminus\{i\}}\frac{|\mathcal{C}|!(|N|-|\mathcal{C}|-1)!}{|N|!}\left(v(\mathcal{C}\cup\{i\})-v(\mathcal{C})\right).$$

Using definition, the sum is over $2^{|N|-1}$ instead of $|N|!$.

Notion of value

Definition (value function)

Let $\mathcal{G}_N$ the set of all TU games (N,v). A **value function** ϕ is a function that assigns to each TU game (N,v) an efficient allocation, i.e. $\phi:\mathcal{G}_N\to\mathbb{R}^{|N|}$ such that $\phi(N,v)(N)=v(N)$.

- The Shapley value is a value function.
- None of the concepts presented thus far were a value function (the nucleolus is guaranteed to be non-empty only for games with a non-empty set of imputations).

Some interesting properties

Let (N,v) and (N,u) be TU games and ϕ be a value function.

- **Symmetry or substitution (SYM):** If $\forall(i,j)\in N$, $\forall\mathcal{C}\subseteq N\setminus\{i,j\}$, $v(\mathcal{C}\cup\{i\})=v(\mathcal{C}\cup\{j\})$ then $\phi_i(N,v)=\phi_j(N,v)$
- **Dummy (DUM):** If $\forall\mathcal{C}\subseteq N\setminus\{i\}$ $v(\mathcal{C})=v(\mathcal{C}\cup\{i\})$, then $\phi_i(N,v)=0$.
- **Additivity (ADD):** Let $(N,u+v)$ be a TU game defined by $\forall\mathcal{C}\subseteq N$, $(u+v)(N)=u(N)+v(N)$. $\phi(u+v)=\phi(u)+\phi(v)$.

Theorem

The Shapley value is the unique value function ϕ that satisfies (SYM), (DUM) and (ADD).

Discussion about the axioms

- SYM: it is desirable that two subsitute agents obtain the same value ✔
- DUM: it is desirable that an agent that does not bring anything in the cooperation does not get any value. ✔
- What does the addition of two games mean?
  - \+ if the payoff is interpreted as an expected payoff, ADD is a desirable property.
  - \+ for cost-sharing games, the interpretation is intuitive: the cost for a joint service is the sum of the costs of the separate services.
  - \- there is no interaction between the two games.
  - \- the structure of the game $(N,v+w)$ may induce a behavior that has may be unrelated to the behavior induced by either games (N,v) or (N,w).
- The axioms are independent. If one of the axiom is dropped, it is possible to find a different value function satisfying the remaining two axioms.

Some properties

Note that other axiomatisations are possible.

Theorem

For superadditive games, the Shapley value is an imputation.

Lemma

For convex game, the Shapley value is in the core.

Simple games

Simple Games

Definition (Simple games)

A game (N, v) is a **Simple game** when

the valuation function takes two values

- 1 for a winning coalitions
- 0 for the losing coalitions

v satisfies *unanimity*: $v(N) = 1$

v satisfies *monotonicity*: $S \subseteq T \Rightarrow v(S) \leqslant v(T)$

One can represent the game by stating all the wining coalitions. Thanks to monotonicity, it is sufficient to only write down the minimal winning coalitions defined as follows:

Definition (Minimal winning coalition)

Let (N, v) be a TU game. A coalition $\mathcal{C}$ is a **minimal winning coalition** iff $v(\mathcal{C}) = 1$ and $\forall i \in \mathcal{C}$, $v(\mathcal{C} \setminus \{i\}) = 0$.

Example

$N = \{1, 2, 3, 4\}$.

We use majority voting, and in case of a tie, the decision of player 1 wins.

The set of winning coalitions is
$\{\{1,2\},\{1,3\},\{1,4\},\{1,2,3\},\{1,2,4\},\{1,3,4\},\{2,3,4\},\{1,2,3,4\}\}$.

The set of minimal winning coalitions is
$\{\{1,2\},\{1,3\},\{1,4\},\{2,3,4\}\}$.

Formal definition of common terms in voting

Definition (Dictator)

Let (N, v) be a simple game. A player $i \in N$ is a **dictator** iff $\{i\}$ is a winning coalition.

Note that with the requirements of simple games, it is possible to have more than one dictator!

Definition (Veto Player)

Let (N, v) be a simple game. A player $i \in N$ is a **veto** player if $N \setminus \{i\}$ is a losing coalition. Alternatively, i is a **veto** player iff for all winning coalition $\mathcal{C}$, $i \in \mathcal{C}$.

It also follows that a veto player is member of every minimal winning coalitions.

Definition (blocking coalition)

A coalition $\mathcal{C} \subseteq N$ is a **blocking coalition** iff $\mathcal{C}$ is a losing coalition and $\forall S \subseteq N \setminus \mathcal{C}$, $S \setminus \mathcal{C}$ is a losing coalition.

Definition (weighted voting games)

A game $(N, w_{i \in N}, q)$ is a **weighted voting game** when v satisfies unanimity, monotonicity and the valuation function is defined as

$$v(S) = \begin{cases} 1 \text{ when } \sum_{i \in S} w_i \geqslant q \\ 0 \text{ otherwise} \end{cases}$$

Unanimity requires that $\sum_{i \in N} w_i \geqslant q$.
If we assume that $\forall i \in N\ w_i \geqslant 0$, monotonicity is guaranteed. For the rest of the lecture, we will assume $w_i \geqslant 0$.

We will note a weighted voting game $(N, w_{i \in N}, q)$ as $[q; w_1, \ldots, w_n]$.

A weighted voting game is a **succint** representation, as we only need to define a weight for each agent and a threshold.

Weighted voting game is a strict subclass of voting games.
i.e., all voting games are **not** weighted voting games.

Examples

- Let us consider the game $[10; 7, 4, 3, 3, 1]$.

  The set of minimal winning coalitions is
  $\{\{1,2\}\{1,3\}\{1,4\}\{2,3,4\}\}$

  Player 5, although it has some weight, is a dummy.

  Player 2 has a higher weight than player 3 and 4, but it is clear that player 2, 3 and 4 have the same influence.
- Let us consider the game $[51; 49, 49, 2]$

  The set of winning coalition is $\{\{1,2\},\{1,3\},\{2,3\}\}$.

  It seems that the players have symmetric roles, but it is not reflected in their weights.

Theorem

Let (N,v) be a simple game. Then

$$Core(N,v) = \left\{ x \in \mathbb{R}^n \;\middle|\; \begin{array}{l} x \text{ is an imputation} \\ x_i = 0 \text{ for each non-veto player } i \end{array} \right\}$$

Proof

- $\subseteq$ Let $x \in Core(N,v)$. By definition $x(N) = 1$. Let i be a non-veto player. $x(N \setminus \{i\}) \geqslant v(N \setminus \{i\}) = 1$. Hence $x(N \setminus \{i\}) = 1$ and $x_i = 0$.
- $\supseteq$ Let x be an imputation and $x_i = 0$ for every non-veto player i. Since $x(N) = 1$, the set V of veto players is non-empty and $x(V) = 1$.

  Let $\mathcal{C} \subseteq N$. If $\mathcal{C}$ is a winning coalition then $V \subseteq \mathcal{C}$, hence $x(\mathcal{C}) \geqslant v(\mathcal{C})$. Otherwise, $v(\mathcal{C})$ is a losing coalition (which may contain veto players), and $x(\mathcal{C}) \geqslant v(\mathcal{C})$. Hence, x is group rational.

□

Shapley-Shubik power index

Definition (Pivotal or swing player)

Let (N,v) be a simple game. A agent i is **pivotal** or **a swing agent** for a coalition $\mathcal{C} \subseteq N \setminus \{i\}$ if agent i turns the coalition $\mathcal{C}$ from a losing to a winning coalition by joining $\mathcal{C}$, i.e., $v(\mathcal{C}) = 0$ and $v(\mathcal{C} \cup \{i\}) = 1$.

Given a **permutation** σ on N, there is a single pivotal agent.

The Shapley-Shubik index of an agent i is the percentage of permutation in which i is pivotal, i.e.

$$I_{SS}(N,v,i) = \sum_{\mathcal{C} \subseteq N \setminus \{i\}} \frac{|\mathcal{C}|!(|N| - |C| - 1)!}{|N|!} \left(v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})\right).$$

"For each permutation, the pivotal player gets a point."

The Shapley-Shubik power index is the Shapley value.
The index corresponds to the expected marginal utility assuming all join orders to form the grand coalitions are equally likely.

Banzhaff power index

Let (N,v) be a TU game.

- We want to count the **number of coalitions** in which an agent is **a swing agent**.
- For each coalition, we determine which agent is a swing agent (more than one agent may be pivotal).
- The **raw Banzhaff index** of a player i is $\beta_i = \dfrac{\sum_{\mathcal{C} \subseteq N \setminus \{i\}} v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})}{2^{n-1}}$.
- For a simple game (N,v), $v(N) = 1$ and $v(\emptyset) = 0$, at least one player i has a power index $\beta_i \neq 0$. Hence, $B = \sum_{j \in N} \beta_j > 0$.
- The **normalized Banzhaff index** of player i for a simple game (N,v) is defined as $I_B(N,v,i) = \dfrac{\beta_i}{B}$.

The index corresponds to the expected marginal utility assuming all coalitions are equally likely.

Examples: [7; 4, 3, 2, 1]

{1,2,3,4}
{1,2,4,3}
{1,3,2,4}
{1,3,4,2}
{1,4,2,3}
{1,4,3,2}
{2,1,3,4}
{2,1,4,3}
{2,3,1,4}
{2,3,4,1}
{2,4,1,3}
{2,4,3,1}
{3,1,2,4}
{3,1,4,2}
{3,2,1,4}
{3,2,4,1}
{3,4,1,2}
{3,4,2,1}
{4,1,2,3}
{4,1,3,2}
{4,2,1,3}
{4,2,3,1}
{4,3,1,2}
{4,3,2,1}

	1	2	3	4
Sh	$\frac{7}{12}$	$\frac{1}{4}$	$\frac{1}{12}$	$\frac{1}{12}$

winning coalitions:
{1,2}
{1,2,3}
{1,2,4}
{1,3,4}
{1,2,3,4}

	1	2	3	4
β	$\frac{5}{8}$	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{1}{8}$
$I_B(N,v,i)$	$\frac{1}{2}$	$\frac{3}{10}$	$\frac{1}{10}$	$\frac{1}{10}$

The Shapley-Shubik index and Banzhaff index may be different.

Representation and Complexitity issues

Representation by enumeration

- Let us assume we want to write a program for computing a solution concept.
- How do we represent the input of a TU game?
- Straighforward representation by enumeration requires **exponential space**.
- Brute force approach may appear good as complexity is measured in term of the **input size**.
- ⇝ we need **compact** or **succint** representation of coalitional games.
- ⇝ e.g., a representation so that the input size is a polynomial in the number of agents.
- In general, the more succint a representation is, the harder it is to compute, hence we look for a balance between succinctness and tractability.

Weighted graph games

A weighted graph game is a coalitional game defined by an undirected weighted graph $\mathcal{G} = (V, W)$ where V is the set of vertices and $W : V \to V$ is the set of edges weights. For $(i,j) \in V^2$, w_{ij} is the weight of the edge between i and j.

- $N = V$, i.e., each agent is a node in the graph.
- for all $\mathcal{C} \subseteq N$, $v(\mathcal{C}) = \sum_{(i,j)\in\mathcal{C}} w_{ij}$.

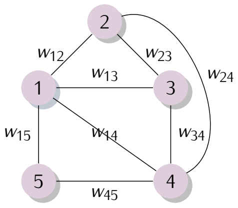

It is a **succint** representation: using an adjacency matrix, we need to provide n^2 entries.
However, it is **not complete**. Some TU games cannot be represented by a weighted graph game (e.g., a majority voting game).

Proposition

Let (V, W) be a weighted graph game. If all the weights are nonnegative then the game is convex.

Proof

$$\begin{aligned} v(S)+v(T) &= \sum_{(i,j)\in S^2} w_{ij} + \sum_{(i,j)\in T^2} w_{ij} = \sum_{(i,j)\in S^2 \vee (i,j)\in T^2} w_{ij} + \sum_{(i,j)\in (S\cap T)^2} w_{ij} \\ &\leqslant \sum_{(i,j)\in (S\cup T)^2} w_{ij} + \sum_{(i,j)\in (S\cap T)^2} w_{ij} = v(S\cup T) + v(S\cap T) \end{aligned}$$

□

Proposition

Let (V, W) be a weighted graph game. If all the weights are nonnegative then membership of a payoff vector in the core can be tested in polynomial time.

Theorem

Let (V, W) a weighted graph game. The Shapley value of an agent $i \in V$ is $Sh_i(N, v) = \frac{1}{2} \sum_{(i,j)\in N^2 \mid i\neq j} w_{ij}$.
The Shapley value can be computed in $O(n^2)$ time.

Proof

Let (V, W) a weighted graph game. We can view this game as the sum of the following $|W|$ games (i.e., one game per edge): $G_{ij} = (V, \{w_{ij}\})$, $(i,j) \in V^2$.
For a game G_{ij}, i and j are substitutes, and all other agents $k \neq i,j$ are dummy agents. Using the symmetry axiom, $Sh_i(G_{ij}) = Sh_j(G_{ij})$. Using the dummy axiom, $Sh_k(G_{ij}) = 0$. Hence, $Sh_i(G_{ij}) = \frac{1}{2} w_{ij}$.
Since (V, W) is the sum of all two-player games, by the additivity axiom, $Sh_k = \sum_{i,j} Sh_k(G_{ij}) = \sum_{k,i} w_{ij}$ □

Theorem

Let (V, W) be a weighted graph game. Testing the nonemptyness of the core is NP-complete.

A representation for superadditive games

Instead of storing a value for each coalition, we can store the positive synergies between agents.
Let (N, v) be a superadditive game and (N, s) its **synergy representation**. Then for a coalition $\mathcal{C} \subseteq N$,

$$v(\mathcal{C}) = \left(\max_{\{\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_k\} \in \mathscr{S}_\mathcal{C}} \sum_{i=1}^{k} v(\mathcal{C}_i) \right),$$

where $\mathscr{S}_\mathcal{C}$ is the set of all partition of $\mathcal{C}$.

Example: $N = \{1,2,3\}$, $v(\{i\}) = 1$, $v(\{1,2\}) = 3$, $v(\{1,3\}) = 2$, $v(\{2,3\}) = 2$, $v(\{1,2,3\}) = 4$.
We can represent this game by $v(\{i\}) = 1$, $v(\{1,2\}) = 3$.

This representation may still require a space exponential in the number of agents, but for many games, the space required is much less.

Theorem

It is NP-complete to determine the value of some coalitions for a coalitional game specified with the synergy representation. In particular, it is NP-complete to determine the value of the grand coalition.

Theorem

Let (N, v) a TU game specified with the synergy representation and the value of the grand coalition.
Then we can determine in polynomial time whether the core of the game is empty.

V. Conitzer and T. Sandholm, **Complexity of constructing solutions in the core based on synergies among coalitions**, *Artificial Intelligence*, 2006.

Multi-issue representation

Some coalitions may form to solve problems requiring distinct competences. For example, solving a set of tasks requiring different expertises.

Definition (Decomposition)

The vector of characteristic functions $\langle v_1, v_2, \ldots, v_T \rangle$, with each $v_i : 2^N \to \mathbb{R}$, is a **decomposition** over T issues of characteristic function $v : 2^N \to \mathbb{R}$
if for any $S \subseteq N$, $v(S) = \sum_{i=1}^{T} v_i(S)$.

It is a fully expressive representation (can use 1 issue).

Multi-issue representation

Theorem
The Shapley value of a coalitional game represented with multi-issue representation can be computed in linear time.

Theorem
Checking whether a given value division is in the core is coNP-complete.

V. Conitzer and T. Sandholm. **Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains**. In *Proc. of the 19th Nat. Conf. on Artificial Intelligence (AAAI-04)*

A logical approach: Marginal contribution nets (MC-nets)

The idea:

- represent each player by a boolean variable,
- treat the characteristic vector of a coalition as a truth assignment.
- the truth assignment can be used to check whether a formula is satisfied and to compute the value of a coalition.

Let N be a collection of atomic variables.

Definition (Rule)
A **rule** has a syntactic form (ϕ, w) where ϕ is called the pattern and is a boolean formula containing variables in N and w is called the weight, and is a real number.

examples:
$(a \wedge b, 5)$: each coalition containing both agents a and b increase its value by 5 units.
$(b, 2)$: each coalition containing b increase its value by 2.

A logical approach: Marginal contribution nets (MC-nets)

Definition (Marginal contribution nets)
An MC-net consists of a set of rules $\{(p_1, w_1), \dots (pk, w_k)\}$ where the valuation function is given by

$$v(\mathcal{C}) = \sum_{i=1}^{k} p_i(e^{\mathcal{C}}) w_i,$$

where $p_i(e^{\mathcal{C}})$ evaluates to 1 if the boolean formula p_i evaluates to true for the truth assignment $e^{\mathcal{C}}$ and 0 otherwise.

S. Ieong and Y. Shoham, **Marginal contribution nets: a compact representation scheme for coalitional games**, in *Proceedings of the 6th ACM conference on Electronic commerce*, 2005.

Examples

Let us consider an MC-net with the following two rules:

$$(a \wedge b, 5) \text{ and } (b, 2)$$

The coalitional game represented has two agents a and b and the valuation function is defined as follows:

$$v(\emptyset) = 0 \qquad v(\{b\}) = 2$$
$$v(\{a\}) = 0 \qquad v(\{a,b\}) = 5+2 = 7$$

We can use negations in rules, and negative weights. Let consider the following example:

$$(a \wedge b, 5),\ (b, 2),\ (c, 4) \text{ and } (b \wedge \neg c, -2)$$

$$v(\emptyset) = 0 \qquad v(\{b\}) = 2-2 = 0 \qquad v(\{a,c\}) = 4$$
$$v(\{a\}) = 0 \qquad v(\{a,b\}) = 5+2-2 = 5 \qquad v(\{b,c\}) = 4+2 = 6$$

Theorem (Expressivity)

- MC-nets can represent **any game** when negative literals are allowed in the patterns, or when the weights can be negative.
- When the patterns are limited to conjunctive formula over postivie literals and weights are nonnegative, MC-nets can represent all and only **convex games**.

Proposition
MC-nets generalize Weighted Graph game representation (strict generalization) and the multi-issue representation.

Theorem
Given a TU game represented by an MC-net limited to conjunctive patterns, the **Shapley value** can be computed in time **linear** in the size of the input.

Proof sketch: we can treat each rule as a game, compute the Shapley value for the rule, and use ADD to sum all the values for the overall game. For a rule, we cannot distinguish the contribution of each agent, by SYM, they must have the same value. It is a bit more complicated when negation occurs (see Ieong and Shoham, 2005).

Proposition
Determining whether the core is empty or checking whether an imputation lies in the core are coNP-hard.

Proof sketch: due to the fact that MC-nets generalize over weighted graph games.

Hedonic Games and NTU games

Hedonic games

Agents have preferences over coalitions, i.e. agent only cares about the other members of the coalition: "enjoying the pleasure of each other's company".

Let N be a set of agents and $\mathcal{N}_i$ be the set of coalitions that contain agent i, i.e., $\mathcal{N}_i = \{\mathcal{C} \cup \{i\} \mid \mathcal{C} \subseteq N \setminus \{i\}\}$.

Definition (Hedonic games)

An **Hedonic game** is a tuple $(N, (\succeq_i)_{i \in N})$ where

- N is the set of agents
- $\succeq_i \subseteq 2^{\mathcal{N}_i} \times 2^{\mathcal{N}_i}$ is a complete, reflexive and transitive preference relation for agent i, with the interpretation that if $S \succeq_i T$, agent i prefers coalition T at most as much as coalition S.

A. Bogomolnaia and M.O. Jackson, *The stability of hedonic coalition structure.* **Games and Economic Behavior**, 2002.

Stability concepts of Hedonic Games

Let $\Pi \in \mathscr{S}_N$ be a coalition structure, and Π_i denotes the coalition in Π that contains i.

- A coalition structure Π is **core stable** iff $\nexists C \subseteq N \mid \forall i \in C, C \succ_i \Pi_i$.
- A coalition structure s is **Nash stable** $(\forall i \in N)\ (\forall C \in \Pi \cup \{\emptyset\})\ \Pi_i \succsim_i C \cup \{i\}$.
  No player would like to join any other coalition in Π assuming the other coalitions did not change.
- A coalition structure Π is **individually stable** iff $\nexists i \in N\ \nexists C \in \Pi \cup \emptyset \mid \big((C \cup \{i\} \succ_i \Pi_i) \wedge (\forall j \in C, C \cup \{i\} \succsim_j C)\big)$.
  No player can move to another coalition that it prefers without making some members of that coalition unhappy.
- A coalition structure Π is **contractually individually stable** iff $\nexists i \in N\ \nexists C \subseteq N \mid$ $(C \cup \{i\} \succ_i \Pi_i) \wedge (\forall j \in C, C \cup \{i\} \succsim_j C) \wedge (\forall j \in \Pi_i \setminus \{i\}, \Pi_i \setminus \{i\} \succsim_j \Pi_i)$
  No player can move to a coalition it prefers so that the members of the coalition it leaves and it joins are better off

Example

$$\{1,2\} \succ_1 \{1\} \succ_1 \{1,2,3\} \succ_1 \{1,3\}$$
$$\{1,2\} \succ_2 \{2\} \succ_2 \{1,2,3\} \succ_2 \{2,3\}$$
$$\{1,2,3\} \succ_3 \{2,3\} \succ_3 \{1,3\} \succ_3 \{3\}$$

{{1,2},{3}} is in the core and is individually stable.
There is no Nash stable partitions.

{{1},{2},{3}}	{1,2} is preferred by both agent 1 and 2, hence not NS, not IS.
{{1,2},{3}}	{1,2,3} is preferred by agent 3, so it is not NS, as agents 1 and 3 are worse off, it is not a possible move for IS. no other move is possible for IS.
{{1,3},{2}}	agent 1 prefers to be on its own (not NS, then, not IS).
{{2,3},{1}}	agent 2 prefers to join agent 1, and agent 1 is better off, hence not NS, not IS.
{{1,2,3}}	agents 1 and 2 have an incentive to form a singleton, hence not NS, not IS.

Nash stability $\Rightarrow$ Individual stability $\Rightarrow$ contractual individual stability

Core stability $\not\Rightarrow$ Nash stability $\not\Rightarrow$ Core stability

Core stability $\not\Rightarrow$ Individual stability

Some classes of games have a non-empty core,
other classes have Nash stable coalition structures.

A. Bogomolnaia and M.O. Jackson, *The stability of hedonic coalition structure.* **Games and Economic Behavior**, 2002.

A representation for hedonic games have been proposed, and is based on MC-nets.

E. Elkind and M. Wooldridge, **Hedonic Coalition Nets**, in *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009

A general model for NTU games (Non-transferable utility games)

It is **not** always possible to compare the utility of two agents or to transfer utility.

Definition (NTU game)

A NTU game is a tuple $(N, X, V, (\succeq_i)_{i \in N})$ where

- X set of outcomes
- $\succeq_i$ a preference relation (transitive and complete) for agent i over the set of outcomes.
- $V(\mathcal{C})$ a set of outcomes that a coalition $\mathcal{C}$ can bring about

- **Example 1:** hedonic games as a special class of NTU games.
  Let $(N, (\succeq_i^H)_{i\in N})$ be a hedonic game.
  • For each coalition $\mathcal{C} \subseteq N$, create a unique outcome $x_{\mathcal{C}}$.
  • For any two outcomes x_S and x_T corresponding to coalitions S and T that contains agent i, We define $\succeq_i$ as follows: $x_S \succeq_i x_T$ iff $S \succeq_i^H T$.
  • For each coalition $\mathcal{C} \subseteq N$, we define $V(\mathcal{C})$ as $V(\mathcal{C}) = \{x_C\}$
- **Example 2:** a TU game can be viewed as an NTU game.
  Let (N,v) be a TU game.
  • We define X to be the set of all allocations, i.e., $X = \mathbb{R}^n$.
  • For any two allocations $(x,y) \in X^2$, we define $\succeq_i$ as follows: $x \succeq_i y$ iff $x_i \geqslant y_i$.
  • For each coalition $\mathcal{C} \subseteq N$, we define $V(\mathcal{C})$ as $V(\mathcal{C}) = \{x \in \mathbb{R}^n \mid \sum_{i\in N} x_i \leqslant v(\mathcal{C})\}$. $V(\mathcal{C})$ lists all the feasible allocation for the coalition $\mathcal{C}$.

Core

An outcome $x \in X$ is **blocked** by a coalition $\mathcal{C}$ if there is some outcome $y \in V(\mathcal{C})$ such that all members i of $\mathcal{C}$ strictly prefer y to x, i.e., $\exists \mathcal{C} \subseteq N$, $\exists y \in V(\mathcal{C})$ s.t. $\forall i \in \mathcal{C}$, $y \succ_i x$.

The **core** of an NTU game $(N, X, V, (\succeq_i)_{i\in N})$ is defined as:
$Core(N,X,V,(\succeq)) = \{x \in V(N) \mid \nexists \mathcal{C} \subseteq N, \nexists y \in V(\mathcal{C}), \forall i \in \mathcal{C} : y \succ_i x\}$

Games with externalities

One of the purpose of Game theory is to *"determine everything that can be said about coalitions between players, compensations between partners in every coalition, mergers or fights between coalitions "...*

von Neumann and Morgenstern,
Theory of games and economic behaviour, 1944.

1- Which coalition will be formed?
2- How will the coalitional worth be shared between members?
3- How does the presence of other coalitions affect the incentives to cooperate?

Cooperative game theory has focused mainly on point 2.

Coalitional Games with externalities

- In a TU game (N,v), the valuation of a coalition depends only on the members, **not** on the other coalition present in the population.
- The value **can** depend on the other coalitions in the population
  - competitive firms
  - teams in sport
- ➪ valuation function for a coalition given a coalition structure (in a competitive setting) $v : 2^N \times \mathscr{S} \to \mathbb{R}$
  Games **in partition function form**.
- ➪ valuation function for each agent given a coalition structure (ex: competitive supply chains) $v : N \times \mathscr{S} \to \mathbb{R}$.
  Games with **Valuations**.

Games in partition function form

Definition (Positive and negative spillovers)
A partition function v exhibits
- **positive spillovers** if for any partition π and any two coalitions S and T in π
  $v(\mathcal{C}, \pi \setminus \{S,T\} \cup \{S \cup T\}) \geqslant v(C,\pi)$ for all coalitions $\mathcal{C} \neq S, T$ in π.
- **negative spillovers** if for any partition π and any two coalitions S and T in π
  $v(\mathcal{C}, \pi \setminus \{S,T\} \cup \{S \cup T\}) \leqslant v(C,\pi)$ for all coalitions $\mathcal{C} \neq S, T$ in π.

Valuations

Assumption: Fixed rules of division appear naturally in many economic situations and in theoretical studies based on a two-stage procedure:

1- formation of the coalitions
2- payoff distribution

Definition (Valuation)

A **valuation** v is a mapping which associates to each coalition structure a payoff of individual payoff in $\mathbb{R}^n$.

Definition (Positive and negative spillovers)

A valuation v exhibits

- **positive spillovers** if for any partition π and any two coalitions S and T in π $v_i(\pi \setminus \{S,T\} \cup \{S \cup T\}) \geqslant v_i(\pi)$ for all players $i \notin S \cup T$.
- **negative spillovers** if for any partition π and any two coalitions S and T in π $v_i(\pi \setminus \{S,T\} \cup \{S \cup T\}) \leqslant v_i(\pi)$ for all players $i \notin S \cup T$.

Definition (Core stability)

A coalition structure π is **core stable** if there does not exist a group $\mathcal{C}$ of players a coalition structure π' that contains $\mathcal{C}$ such that $\forall i \in \mathcal{C}$, $v_i(\pi') > v_i(\pi)$.

Definition (α-core Stability)

A coalition structure π is **α-core stable** if there does not exist a group $\mathcal{C}$ of players and a partition $\pi'_{\mathcal{C}}$ such that, for all partition $\pi_{N \setminus \mathcal{C}}$ formed by external players, $\forall i \in \mathcal{C}$, $v_i(\pi'_{\mathcal{C}} \cup \pi_{N \setminus \mathcal{C}}) > v_i(\pi)$.

Definition (β-core Stability)

A coalition structure π is **β-core stable** if there does not exist a group $\mathcal{C}$ of players such that for all partitions $\pi_{N \setminus \mathcal{C}}$ of external players, there exists a partition $\pi_{\mathcal{C}}$ of $\mathcal{C}$ such that $\forall i \in \mathcal{C}$, $v_i(\pi_{\mathcal{C}} \cup \pi_{N \setminus \mathcal{C}}) > v_i(\pi)$.

Issues studied in multiagent systems

Search of an optimal coalition structure

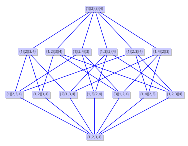

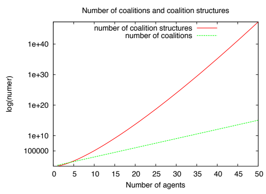

The difficulty of searching for the optimal CS is the large search space.

How to distribute the computation of all the coalition values?

- goal is to minimize computational time
  Computing the value of a coalition can be hard: ex solving a TSP
- load balancing: distribute coalitions of every size equally among the agents coalitions.

but agents may have different compuatational speed

A naive approach does not avoid redundancy and may have a high communication complexity.
The current best algorithm works by sharing the computation of coalition of the same size between all the agents.

O. Shehory and S. Kraus. **Methods for task allocation via agent coalition formation**. *Artificial Intelligence*, 1998

T. Rahwan and N. Jennings. **An algorithm for distributing coalitional value calculations among cooperating agents**, *Artificial Intelligence*, 2007

Search of the Optimal Coalition Structure

- First algorithm that guarantees a bound from the optimal $\frac{v(s)}{V(s^\star)} \leqslant \mathcal{K}$. It is necessary to visit a least 2^{n-1} CSs, which corresponds to the first two levels of the lattice.
- Best current algorithm is called IP for Interger Partition:
  - Integer Partition: ex $[1,1,2] \to$ space of coalition structures containing two singletons and a coalition of size 2.
  - Finding bounds for each subspace is easy. Ex:
    $\max_{S \in [1,1,2]} v(S) \leqslant \max_{\mathcal{C} \in 2^N, |\mathcal{C}|=1} v(\mathcal{C}) + \max_{\mathcal{C} \in 2^N, |\mathcal{C}|=1} v(\mathcal{C}) + \max_{\mathcal{C} \in 2^N, |\mathcal{C}|=2} v(\mathcal{C})$
  - IP uses the representation to efficiently prune part of the space and search the most promising subspaces.

T. Rahwan, S.D. Ramchurn, N. Jennings, and A. Giovannucci. **An anytime algorithm for optimal coalition structure generation**, *Journal of Artificial Intelligence Research*, 2009.

Environments, Safety and Robustness, Communication

- Agents can enter and leave the environment at any time
- The characteristics of the agents may change with time
- Communication links may fail during the negotiation phase

Extending some concepts to Open Environments.

⇝ how to avoid recomputing from scratch?

Additional goals of the coalition formation: decreasing the time and the number of messages required to reach an agreement.

- ➪ learning may be used to decrease negotiation time.
- ➪ communication costs are represented in the characteristic function.
- ➪ analysis of the communication complexity of computing the payoff of a player with different stability concepts: they find that it is $\Theta(n)$ when the Shapley value, the nucleolus, or the core is used.

Uncertainty about Knowledge and Task

- Agents may not know some tasks.
- Agents may not know the valuation function, and may use Fuzzy sets to represent the coalition value.
- Expected values of coalitions are used instead of the valuation function.
- Approximation of valuation function: e.g., computing a value for a coalition requires solving a version of the traveling salesman problem and approximations are used to solve that problem.
- Agent do not know the cost incurred by other agents and may only estimate these costs.

Manipulation

A protocol may require that they disclose some private information.

- ➪ Avoid information asymmetry that can be exploited by some agents by using cryptographic techniques.
- ➪ Use computational complexity to protect a protocol.

Other types of manipulations:

- hiding skills
- using false names (anonymous environments)
- colluding

The traditional solution concepts can be vulnerable to false names and to collusion.
Study for some TU games and for weighted voting games.

Long Term Vs Short Term

In general, a coalition is a short-lived entity that is *"formed with a purpose in mind and dissolve when that need no longer exists, the coalition ceases to suit its designed purpose, or critical mass is lost as agents depart"*.

- Long term coalitions, and in particular the importance of trust in this content.
- Repeated coalition formation under uncertainty using learning.

Overlapping Coalitions

Agents may simultaneously belong to more than one coalition

- ➪ Fuzzy approach
  - agents can be member of a coalition with a certain degree that represents the risk associated with being in that coalition.
  - agents have different degree of membership, and their payoff depends on this degree.
- ➪ Heuristic algorithms
- ➪ Game theoretical approach (overlapping core)

Conclusion

- Game theory proposes many solution concepts (some of which were not introduced: bargaining set, ϵ-core, least-core, Owen value). Each solution concept has pros and cons.
- Work in AI has dealt with representation issues, and practical coalition formation protocols.
- Many issues are left unexaplored.

Lecture Notes on Cooperative Games in Multiagent Systems*

Stéphane Airiau
s.airiau@uva.nl

Institute for Logic, Language & Computation
Universiteit van Amsterdam, The Netherlands

Contents

*Notes prepared for a tutorial at the 12^{th} European Agent Systems Summer School (EASSS-2010), Saint Etienne, France, 23rd-24th August 2010.

1 Introduction

Coalition formation is an important tool for enabling cooperation in agent societies. Social scientists and economists have studied situations where individuals and businesses benefit from joining forces. The coalition formation problem can be decomposed into two sub-problems. The first problem is the *selection* of the coalition members in the agent population. Then, it is assumed that the members self-organise and achieve their goals, and that the coalition *as a whole* receives a value, i.e., the cooperation of the coalition members is rewarded, not the individual agents. The second problem is the *sharing* of this value between its members. These two sub-problems cannot be treated independently [76]: a rational agent will not accept to be a member of a coalition if it can get a higher reward by joining a different coalition. Games that model such cooperation have been extensively studied in the game theory literature in a sub-field called “Cooperative game theory”. A central point of attention is the stability of the coalition: it is preferable that agents do not have any incentive to leave their current coalition to join a different one. Unfortunately, there is no unique and accepted solution to enforce stability, there are different stability criteria, with their own strengths and weaknesses.

Over the last decade, coalition formation has received increased attention in the multiagent systems community: forming dynamic coalitions may lead to more efficient artificial agent societies. Joining a coalition may be beneficial for an agent: the use of other members’ resources may facilitate or enable the solution of a problem. Cooperative game theory has provided a great basis to build coalition formation protocols, but additional issues have risen while trying to apply them.

Multiple scenarios, for example in the task allocation domains [82] or in the electronic marketplace domain [89], have inspired many coalition formation protocols, many of them based on game-theoretic stability criteria, which guarantees a fair distribution of the value among the coalition members. These scenarios have also brought to light many issues and constraints that classical game theory did not address. For example, one issue is the complexity of computing the stability criteria. Communication complexity can also limit the use of game-theoretic concepts. Some other issues are related with dynamic environments: agents can enter and leave the system at any time, new tasks may appear in the environment, the environment may be uncertain (uncertainty about the value of the coalitions, about the competence of other agents, etc.). Safety and robustness issues should also be taken into account to guarantee a stable agent society. In addition, researchers must design protocols that are secure to prevent the possibility of manipulation or infiltration by agents or external forces. Another scenario is to consider that the goal of the agent is to maximise utilitarian social welfare. This scenario is not interesting for game theory as sharing the value between the members is no longer an issue. However, finding the optimal organisation is still a hard problem which can be addressed by AI techniques.

Section 2 and 3 survey the field of cooperative game theory. Section 2 treats the case in which utility can be transferred between agents: the transferable utility games (TU games). This assumes that 1) interpersonal comparison of utility is possible and 2) it is possible to transfer utility between agents. We introduce various solution concepts for TU games (some that focus on the stability of the coalition (core, nucleolus and kernel) and one that focuses on fairness (the Shapley value). We then introduce a special class of TU games that can model voting situation, and in particular, we will introduce measures of the voting power of the agents. In Section 3, we will introduce the case where no transfer or comparison of utility are possible between agents: the non-transferable utility games (NTU games). We will first introduce a special class of games called hedonic games to introduce the concept. Then, we provide the general definition of an NTU game and some definitions of stability concepts.

The rest of the paper introduces research in coalition formation in the multiagent systems literature. In Section 4, we consider domains where the agents’ goal is to maximise utilitarian social welfare: i.e., the case agents’ goal is to maximize the utility of the entire society of agents. We survey some central algorithms that efficiently search for the optimal partition of agents into coalitions. Then in Section 5, we present some applications that have motivated many study in multiagent systems. In In particular, we discuss the task allocation domain, the electronic marketplace domain, and some variants. We also list some additional domains where coalitions of agents have been used. Finally, in Section 6, we survey the issues raised by the multiagent systems community for the problem of coalition formation for which game theory has little (or no) answer so far.

2 TU games

The game theory community has extensively studied the coalition formation problem [41, 61]. The literature is divided into two main models, depending on whether utility can be transferred between individuals. In a transferable utility game (or TU game), it is assumed that agents can compare their utility and that a common scale of utility exists. In this case, it is possible to define a value for a coalition as the worth the coalition can achieve through cooperation. The agents have to share the value of the coalition, hence utility needs to be transferable. In a so-called non-transferable utility game (or NTU game), inter-personal comparison of utility is not possible, and agents have a preference over the different coalitions of which it is a member. In this section, we introduce the TU games.

2.1 Definitions

In the following, we use a utility-based approach and we assume that "everything has a price": each agent has a utility function that is expressed in currency units. The use of a common currency enables the agents to directly compare alternative outcomes, and it also enables side payments. The definition of a TU game is simple: it involves a set of players and a characteristic function (a map from sets of agents to real numbers) which represents the value that a coalition can achieve. The characteristic function is common knowledge and the value of a coalition depends only on the other players present in its coalition.

2.1.1 Notations and types of TU games

We consider a set N of n agents. A *coalition* is a non-empty subset of N. The set N is also known as the *grand coalition*. The set of all coalitions is 2^N and its cardinality is 2^n. A *coalition structure (CS)* $\mathcal{S} = \{\mathcal{C}_1, \cdots, \mathcal{C}_m\}$ is a partition of N: each set $\mathcal{C}_i$ is a coalition with $\cup_{i=1}^{m} \mathcal{C}_i = N$ and $i \neq j \Rightarrow \mathcal{C}_i \cap \mathcal{C}_j = \emptyset$. We will denote $\mathscr{S}_\mathcal{C}$ the set of all partitions of a set of agnets $\mathcal{C} \subseteq N$. The set of all CSs is then denoted as $\mathscr{S}_N$, its size is of the order $O(n^n)$ and $\omega(n^{\frac{n}{2}})$ [76]. The *characteristic function* (or *valuation function*) $v : 2^N \to \mathbb{R}$ provides the worth or utility of a coalition. Note that this definition assumes that the valuation of a coalition $\mathcal{C}$ does not depend on the other coalitions present in the population.

Definition 2.1 (TU game)**.** *A transferable utility game (TU game) is defined as a pair (N, v) where N is the set of agents, and $v : 2^N \to \mathbb{R}$ is a characteristic function.*

A first example of a TU game is the *majority game*. Assume that the number of agents n is odd and that the agents decide between two alternatives using a majority vote. Also assume that no agent is indifferent, i.e., an agent always strictly prefers one alternative over the other. We model this by assigning to a "winning coalition" the value 1 and to the other ones the value 0, i.e.,

$$v(\mathcal{C}) = \begin{cases} 1 \text{ when } |\mathcal{C}| > \frac{n}{2} \\ 0 \text{ otherwise} \end{cases}$$

We now describes some types of valuation functions.

Additive (or inessential): $\forall \mathcal{C}_1, \mathcal{C}_2 \subseteq N \mid \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$, $v(\mathcal{C}_1 \cup \mathcal{C}_2) = v(\mathcal{C}_1) + v(\mathcal{C}_2)$. When a TU game is additive, $v(\mathcal{C}) = \sum_{i \in \mathcal{C}} v(i)$, i.e., the worth of each coalition is the same whether its members cooperate or not: there is no gain in cooperation or any synergies between coalitions, which explains the alternative name (inessential) used for such games.

Superadditive: $\forall \mathcal{C}_1, \mathcal{C}_2 \subseteq N \mid \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$, $v(\mathcal{C}_1 \cup \mathcal{C}_2) \geq v(\mathcal{C}_1) + v(\mathcal{C}_2)$, in other words, any pair of coalitions is best off by merging into one. In such environments, social welfare is maximised by forming the grand coalition.

Subadditive: $\forall \mathcal{C}_1, \mathcal{C}_2 \subseteq N \mid \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$, $v(\mathcal{C}_1 \cup \mathcal{C}_2) \leq v(\mathcal{C}_1) + v(\mathcal{C}_2)$: the agents are best off when they are on their own, i.e., cooperation not desirable.

Convex games: First let us call $v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})$ the marginal contribution of a player i to coalition $\mathcal{C}$, i.e., it is the increase of value of coalition $\mathcal{C}$ due to the presence of agent i. We call then a valuation convex if for all $\mathcal{C} \subseteq T$ and $i \notin T$ $v(\mathcal{C} \cup \{i\}) - v(\mathcal{C}) \leq v(T \cup \{i\}) - v(T)$. So a valuation function is convex when the marginal contribution of each player increases with the size of the coalition he joins. Convex valuation functions are superadditive.

Monotonic A function is monotonic when $\forall \mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq N$, $v(C_1) \leq v(\mathcal{C}_2)$. In other words, when more agents join a coalition, the value of the larger coalition is at least the value of the smaller one. For example, the valuation function of the majority game is monotonic: when more agents join a coalition, they cannot turn the coalition from a winning to a losing one.

Unconstrained. The valuation function can be superadditive for some coalitions, and subadditive for others: some coalitions should merge when others should remain separated. This is the most difficult and interesting environment.

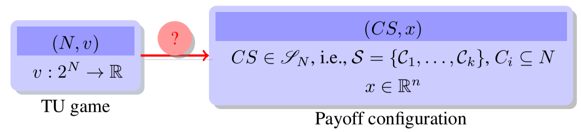

Figure 1: What is solving TU games?

The valuation function provides a value to a set of agents, not to individual agents. The *payoff distribution* $x = \{x_1, \cdots, x_n\}$ describes how the worth of the coalition is shared between the agents, where x_i is the payoff of agent i. We also use the notation $x(\mathcal{C}) = \sum_{i \in \mathcal{C}} x(i)$. A *payoff configuration (PC)* is a pair $(\mathcal{S}, x)$ where $\mathcal{S} \in \mathscr{S}_N$ is a CS and x is a payoff distribution. Given a TU game (N, v) as an input, the *fundamental question* is what PC will form: what are the coalitions that will form and how to distribute the worth of the coalition (see Figure 1).

$$\begin{gathered} N = \{1, 2, 3\} \\ v(\{1\}) = 0, v(\{2\}) = 0, v(\{3\}) = 0 \\ v(\{1, 2\}) = 90 \\ v(\{1, 3\}) = 80 \\ v(\{2, 3\}) = 70 \\ v(\{1, 2, 3\}) = 105 \end{gathered}$$

Table 1: An example of a TU game

Let us go over the TU game in Table 1. In this example, there are three agents named 1, 2 and 3. There are 7 possible coalitions and the value of each coalition is given in the table. There are 5 CSs which are the following: $\{\{1\}, \{2\}, \{3\}\}$, $\{\{1, 2\}, \{3\}\}$, $\{\{1\}, \{2, 3\}\}$, $\{\{2\}, \{1, 3\}\}$, $\{\{1, 2, 3\}\}$. What PC should be chosen? Should the agents form the grand coalition and share equally the value? The choice of the coalition can be justified by arguing that it is the coalition that generates the most utility for the society. However, is an equal share justified? Agent 3 could propose to agent 1 to form $\{1, 3\}$ and to share equally the value of this coalition (hence, 40 for each agent). Actually, agent 2 can make a better offer to agent 1 by proposing an equal share of 45 if they form $\{1, 2\}$. Agent 3 could then propose to agent 1 to form $\{1, 3\}$ and to let it get 46 (agent 3 would then have 34). Is there a PC that would be preferred by all agents at the same time?

2.1.2 Rationality concepts

We now discuss different rationality concepts for payoff distributions, i.e., some desirable properties that link the coalition values to the agents' individual payoff.

Efficiency: $x(N) = v(N)$ the payoff distribution is an allocation of the whole worth of the grand coalition to all the players. In other words, no utility is lost at the level of the population.

Individual rationality: An agent i will be a member of a coalition only when $x_i \geq v(\{i\})$, i.e., to be part of a coalition, a player must be better off than when it is on its own.

Group rationality: $\forall \mathcal{C} \subseteq N$, $x(\mathcal{C}) \geq v(\mathcal{C})$, i.e., the sum of the payoff of a coalition should be at least the value of the coalition (there should not be any loss at the level of a coalition).

Pareto optimal payoff distribution: It may be desirable to have a payoff distribution where no agent can improve its payoff without lowering the payoff of another agent. More formally, a payoff distribution x is Pareto optimal iff $\nexists y \in \mathbb{R}^n \mid \exists i \in N \mid \{y_i > x_i$ and $\forall j \neq i,\, y_j \geq x_j\}$.

Two notions will be helpful to discuss some solution concepts. The first is the notion of *imputation*, which is a payoff distribution with the minimal acceptable constraints.

Definition 2.2 (Imputation)**.** *An* imputation *is a payoff distribution that is efficient and individually rational for all agents.*

An imputation is a solution candidate for a payoff distribution, and can also be used to object a payoff distribution. The second notion is the *excess* which can be seen as an amount of complaint or as a potential strength depending on the view point.

Definition 2.3 (Excess)**.** *The excess related to a coalition* $\mathcal{C}$ *given a payoff distribution* x *is* $e(\mathcal{C}, x) = v(\mathcal{C}) - x(\mathcal{C})$.

When $e(\mathcal{C}, x) > 0$, the excess can be seen as an amount of complaint for the current members of $\mathcal{C}$ as some part of the value of the coalition is lost. When $\mathcal{C}$ is not actually formed, some agent $i \in \mathcal{C}$ can also see the excess as a potential increase of its payoff if $\mathcal{C}$ was to be formed. Some stability concepts (the kernel and the nucleolus, see below) are based on the excess of coalitions. Another stability concept can also be defined in terms of the excess.

2.2 The core

Let us assume that we have a TU game (N, v) and that we want to form the grand coalition. What would be a good payoff distribution x. Following the discussion about the game in Table 1, it seems desirable that no agent has an incentive to form a different coalition such that all members benefit. This is the idea about the core, which was first introduced by Gillies [35]. A payoff distribution is in the *core* when no agent has any incentive to be in a different coalition. This is at first akin to the Nash equilibrium concept of non-cooperative games. The core actually presents a stronger condition: no set of agents can benefit by forming a new coalition, which corresponds to the group rationality assumption. More formally:

Definition 2.4 (Core)**.** *A payoff distribution* $x \in \mathbb{R}^n$ *is in the* core *of a game* (N, v) *iff* x *is an imputation that is group rational, i.e.,* $core(N, v) = \{x \in \mathbb{R}^n \mid \sum_{i \in N} x_i = v(N) \ \wedge\ \forall \mathcal{C} \subseteq N\ x(\mathcal{C}) \geq v(\mathcal{C})\}$

A payoff distribution is in the core when no group of agents has any interest in rejecting it, i.e., no group of agents can gain by forming a different coalition. Note that this condition has to be true for all subsets of N, in particular, this ensures individual rationality. The core can thus be defined as a payoff structure that satisfies weak linear inequalities. The core is therefore closed and convex. Another way to define the core is in terms of excess:

Definition 2.5 (Core)**.** *The core is the set of payoff distribution* $x \in \mathbb{R}^n$, *such that* $\forall R \subset N$, $e(R, x) \leq 0$

In other words, a PC is in the core when there exists no coalition that has a positive excess. This definition is attractive as it shows that no coalition has any complaint: each coalition's demand can be granted.

Let us consider the following two-player game $(\{1, 2\}, v)$ where $v(\{1\}) = 5$, $v(\{2\}) = 5$, and $v(\{1, 2\}) = 20$. The core of the game is a segment defined as follows: $core(N, v) = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 \geq 5,\ x_2 \geq 5,\ x_1 + x_2 = 20\}$ and is represented in Figure 2. This example shows that, although the game is symmetric, most of the payoffs in the core are not fair. Core allocations focus on stability only and they may not be fair.

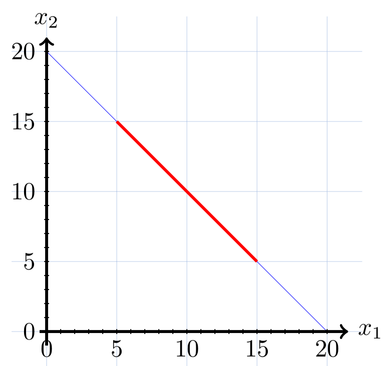

Figure 2: Example of a core allocation

It is possible to represent the core for game with three agents. For a game $(\{1, 2, 3\}, v)$, the efficiency condition is $v(\{1, 2, 3\}) = x_1 + x_2 + x_3$, which is a plane in a 3-dimensional space. On this plane, we can draw the conditions for individual rationality and for group rationality. Each of these conditions partitions the space into two regions separated by a line: one region is incompatible with a core allocation, the other region is. The core is the intersection of the compatible regions. Figure 3 represents the core of a three-player game.

There are, however, multiple concerns associated with using the notion of the core. First, the core can be empty: the conflicts captured by the characteristic function cannot satisfy all the players simultaneously. When

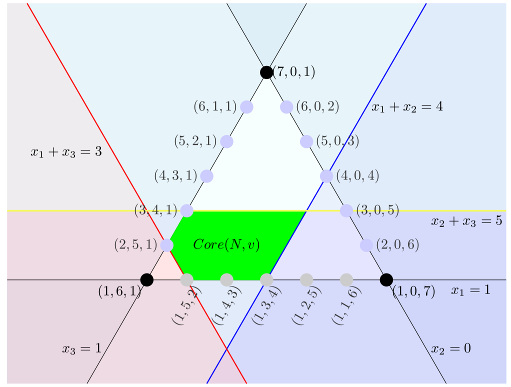

$$v(\{1\}) = 1 \quad v(\{2\}) = 0 \quad v(\{3\}) = 1$$
$$v(\{1,2\}) = 4 \quad v(\{1,3\}) = 3 \quad v(\{2,3\}) = 5$$
$$v(\emptyset) = 0 \quad v(\{1,2,3\}) = 8$$

Figure 3: Example of a three-player game: The core is the area in green

the core is empty, at least one coalition is dissatisfied by the utility allocation and therefore is a blocking coalition. Let us consider the following example from [41]: $v(\{A, B\}) = 90$, $v(\{A, C\}) = 80$, $v(\{B, C\}) = 70$, and $v(N) = 120$. In this case, the core is the PC where the grand coalition forms and the associated payoff distribution is $(50, 40, 30)$. If $v(N)$ is increased, the size of the core also increases. But if $v(N)$ decreases, the core becomes empty (this was the case of the example in Table 1).

exercise: How can you modify the game in Figure 3 so that the core becomes empty?

Some classes of games, however, are guaranteed to have a non-empty core. The following theorem covers the class of *convex games*.

Theorem 1. *A convex game has a non-empty core.*

Proof. Let us assume a convex game (N, v). Let us define a payoff vector x in the following way: $x_1 = v(\{1\})$ and for all $i \in \{2, \ldots, n\}$, $x_i = v(\{1, 2, \ldots, i\}) - v(\{1, 2, \ldots, i-1\})$. In other words, the payoff of the i^{th} agent is its marginal contribution to the coalition consisting of all previous agents in the order $\{1, 2, \ldots, i-1\}$.

Let us prove that the payoff vector is *efficient* by writing up and summing the payoff of all agents:

$$\begin{array}{rcl}
x_1 & = & v(\{1\}) \\
x_2 & = & v(\{1,2\} - v(\{1\}) \\
 & \cdots & \\
x_i & = & v(\{1,2,\ldots,i\}) - v(\{1,2,\ldots,i-1\}) \\
 & \cdots & \\
x_n & = & v(\{1,2,\ldots,n\}) - v(\{1,2,\ldots,n-1\}) \\
\hline
\sum_{i \in N} x_n & = & v(\{1,2,\ldots,n\}) = v(N)
\end{array}$$

By summing these n equalities, we obtain the efficiency condition: $\sum_{i \in N} x_n = v(\{1, 2, \ldots, n\}) = v(N)$.✔

Let us prove that the payoff vector is *individually rational*. By convexity, we have $v(\{i\}) - v(\emptyset) \leq v(\{1, 2, \ldots, i\}) - v(\{1, 2, \ldots, i-1\})$, hence $v(\{i\}) \leq x_i$. ✔

Finally, let us prove that the payoff vector is *group rational*. Let $C \subseteq N$, $C = \{a_1, a_2, \ldots, a_k\}$ and let us consider that $a_1 < a_2 < \ldots < a_k$. It is obvious that $\{a_1, a_2, \ldots, a_k\} \subseteq \{1, 2, \ldots, a_k\}$. Using the convexity assumption, we obtain the following:

$$\begin{array}{rcl}
v(\{a_1\}) - v(\emptyset) & \leq & v(\{1,2,\ldots,a_1\}) - v(\{1,2,\ldots,a_1-1\}) = x_{a_1} \\
v(\{a_1,a_2\}) - v(\{a_1\}) & \leq & v(\{1,2,\ldots,a_2\}) - v(\{1,2,\ldots,a_2-1\}) = x_{a_2} \\
 & \cdots & \\
v(\{a_1,a_2,\ldots,a_l\}) - v(\{a_1,a_2,\ldots,a_{l-1}\}) & \leq & v(\{1,2,\ldots,a_l\}) - v(\{1,2,\ldots,a_l-1\}) = x_{a_l} \\
 & \cdots & \\
v(\{a_1,a_2,\ldots,a_k\}) - v(\{a_1,a_2,\ldots,a_{k-1}\}) & \leq & v(\{1,2,\ldots,a_k\}) - v(\{1,2,\ldots,a_k-1\}) = x_{a_k} \\
\hline
v(C) = v(\{a_1,a_2,\ldots,a_k\}) & \leq & \sum_{i=1}^{k} x_{a_k} = x(C)
\end{array}$$

By summing these k inequalities, we obtain $v(C) = v(\{a_1, a_2, \ldots, a_k\}) \leq \sum_{i=1}^{k} x_{a_k} = x(C)$, which is the group rationality condition. ✔ □

Another example are minimum cost spanning tree game. Let N be the set of customers, and let 0 be the supplier. Let us define $N_* = N \cup \{0\}$. For $(i,j) \in N_*^2$, $i \neq j$, let $c_{i,j}$ be the cost of connecting i and j by the edge e_{ij}. Let (N,c) be the corresponding cost game, which is called a minimum cost spanning tree game ((MCST) game).

Theorem 2. *Every minimum cost spanning tree game has a non-empty core.*

Proof. Let us define a cost distribution x and then we will show that x is in the core.

Let $T = (N, E_N)$ a minimum cost spanning tree for the graph $\left(N_*, c_{\{ij\} \subseteq N_*^2}\right)$. Let i be a customer. Since T is a tree, there is a unique path $(0, a_1, \ldots, a_k, i)$ from 0 to i. The cost paid by agent i is defined by $x_i = c_{a_k,i}$.

This cost allocation is efficient by construction of x.

We need to show the cost allocation is group rational, i.e. for all coalition S, we have $x(S) \leq v(S)$ (it is a cost, which explains the inequality). Let $S \subset N$ and $T_S = (S \cup \{0\}, E_s)$ be a minimum cost spanning tree of the graph $\left(S \cup \{0\}, c_{\{ij\} \in S \cup \{0\}}\right)$. Let extand the tree T_S to a graph $T_S^+ = (N_*, E_N^+)$ by adding the remaining customers $N \setminus S$, and for each customer $i \in N \setminus S$, we add the edge of E_N ending in i, i.e., we add the edge (a_k, i). The graph T_S^+ has $|S| + |N \setminus S|$ edges an is connected. Hence, T_S^+ is a spanning tree. Now, we note that $c(S) + x(N \setminus S) = \sum_{e_{ij} \subseteq E_N^+} c_{ij} \geq \sum_{e_{ij} \subseteq E_N} = c(N) = x(N)$. The inequality is due to the fact that T_S^+ is a spanning tree, and T is a minimum spanning tree. It follows that $x(S) \leq v(S)$. ✔ □

The set of games with non-empty core has been characterized independently by Bondareva (1963) and Shapley (1967), and the result is known as the Bondareva Shapley theorem. This result connects results from linear programming with the concept of the core.

Let $\mathcal{C} \subseteq N$. The *characteristic vector* $\chi_\mathcal{C}$ of $\mathcal{C}$ is the member of $\mathbb{R}^N$ defined by $\chi^i_\mathcal{C} = \begin{cases} 1 \text{ if } i \in \mathcal{C} \\ 0 \text{ if } i \in N \setminus \mathcal{C} \end{cases}$

Definition 2.6 (Map)**.** *A* map *is a function* $2^N \setminus \emptyset \to \mathbb{R}_+$ *that gives a positive weight to each coalition.*

Definition 2.7 (Balanced map)**.** *A function* $\lambda : 2^N \setminus \emptyset \to \mathbb{R}_+$ *is a* balanced map *iff* $\sum_{\mathcal{C} \subseteq N} \lambda(\mathcal{C}) \chi_\mathcal{C} = \chi_N$

A map is balanced when the amount received over all the coalitions containing an agent i sums up to 1. We provide an example in Table 2 for a three-player game.

	1	2	3
$\{1,2\}$	$\frac{1}{2}$	$\frac{1}{2}$	0
$\{1,3\}$	$\frac{1}{2}$	0	$\frac{1}{2}$
$\{2,3\}$	0	$\frac{1}{2}$	$\frac{1}{2}$

$\lambda(\mathcal{C}) = \begin{cases} \frac{1}{2} \text{ if } |\mathcal{C}| = 2 \\ 0 \text{ otherwise} \end{cases}$

Each of the column sums up to 1.

$\frac{1}{2}\chi_{\{1,2\}} + \frac{1}{2}\chi_{\{1,3\}} + \frac{1}{2}\chi_{\{2,3\}} = \chi_{\{1,2,3\}}$

Table 2: Example of a balanced map for $n = 3$

Definition 2.8 (Balanced game)**.** *A game is* balanced *iff for each balanced map λ we have* $\sum_{\mathcal{C} \subseteq N, \mathcal{C} \neq \emptyset} \lambda(\mathcal{C}) v(\mathcal{C}) \leq v(N)$.

Theorem 3 (Bondareva-Shapley theorem)**.** *A TU game has a non-empty core iff it is balanced.*

This theorem completely characterize the set of games with a non-empty core. However, it is not always easy or feasible to check that it is a balanced game.

There are few extensions to the concept of the core. As discussed above, one main issue of the core is that it can be empty. In particular, a member of a coalition may block the formation so as to gain a very small payoff. When the cost of building a coalition is considered, it can be argued that it is not worth blocking a coalition for a small utility gain. The strong and weak ϵ-core concepts model this possibility. The constraints defining the strong (respectively the weak) ϵ-core become $\forall T \subseteq N, x(T) \geq v(T) - \epsilon$, (respectively $\forall T \subseteq N, x(T) \geq v(T) - |T| \cdot \epsilon$). In the weak core, the minimum amount of utility required to block a coalition is per player, whereas for the strong core, it is a fixed amount.

2.3 Games with Coalition Structure

In most traditional work in game theory, the superadditivity of the valuation function is not explicitly stated, but it is implicitly assumed. In the definition of the core that we presented, we stated that the grand coalition is formed. With this definition, checking whether the core is empty amounts to checking whether the grand coalition is stable.

Aumann and Drèze discuss why the coalition formation process may generate a CS that is not the grand coalition [9]. One reason they mention is that the valuation may not be superadditive (and they provide some discussion about why it may be the case). Another reason is that a CS may "reflect considerations that are excluded from the formal description of the game by necessity (impossibility to measure or communicate) or by choice" [9]. For example, the affinities can be based on location (each agent may come from the same country or university), or trust relations, etc.

Definition 2.9 (Game with coalition structure)**.** *A* game with coalition structure *is a triplet* (N, v, S)*, where* (N, v) *is a TU game, and S is a particular CS. In addition, transfer of utility is only permitted within (not between) the coalitions of S, i.e.,* $\forall \mathcal{C} \in S$, $x(\mathcal{C}) = v(\mathcal{C})$.

Another way to understand this definition is to consider that the problems of deciding which coalition forms and how to share the coalition's payoff are decoupled: the choice of the coalition is made first and results in the CS. Only the payoff distribution choice is left open. The agents are allowed to refer to the value of coalition with agents oustide of their coalition (i.e., opportunities they would get outside of their coalition) to negotiate a better payoff. Note that the agents' goal is not to change the CS, but only to negotiate a better payoff for themselves.

Aumann and Drèze extend the definition of many solution concepts and we now present the definition of the core, which is borrowed from [21]. The definition is similar to [31, 75]. First, we need to define the set of possible payoffs: the payoff distributions such that the sum of the payoff of the members of a coalition in the CS does not exceed the value of that coalition. More formally:

Definition 2.10. *Feasible payoff Let* $(N, v, \mathcal{S})$ *be a TU game with CS. The set of* feasible payoff distributions *is* $X_{(N,v,\mathcal{S})} = \{x \in \mathbb{R}^n \mid \forall \mathcal{C} \in \mathcal{S} x(\mathcal{C}) \leq v(\mathcal{C})\}$.

A payoff distribution x is *efficient* with respect to a CS $\mathcal{S}$ when $\forall \mathcal{C} \in \mathcal{S}$, $\sum_{i \in \mathcal{C}} x_j = v(\mathcal{C})$. A payoff distribution is an *imputation* when it is efficient (with respect to the current CS) and individually rational (i.e., $\forall i \in N$, $x_i \geq v(\{i\})$). The set of all imputations for a CS $\mathcal{S}$ is denoted by $\mathcal{I}mp(\mathcal{S})$. We can now state the definition of the core:

Definition 2.11 (Core)**.** *The core of a game* $(N, v, \mathcal{S})$ *is the set of all PCs* $(\mathcal{S}, x)$ *such that* $x \in \mathcal{I}mp(\mathcal{S})$ *and* $\forall \mathcal{C} \subseteq N$, $\sum_{i \in \mathcal{C}} x_j \geq v(\mathcal{C})$*, i.e.,* $core(N, v, \mathcal{S}) = \{x \in \mathbb{R}^n \mid (\forall \mathcal{C} \in \mathcal{S}, x(\mathcal{C}) \leq v(\mathcal{C})) \wedge (\forall \mathcal{C} \subseteq N, x(\mathcal{C}) \geq v(\mathcal{C}))\}$.

We now provide a theorem by Aumann and Drèze which shows that the core satisfies a desirable properties: if two agents can be substituted, then a core allocation must provide them identical payoffs.

Definition 2.12 (Substitutes)**.** *Let* (N, v) *be a game and* $(i, j) \in N^2$*. Agents i and j are* substitutes *iff* $\forall \mathcal{C} \subseteq N \setminus \{i, j\}$, $v(\mathcal{C} \cup \{i\}) = v(\mathcal{C} \cup \{j\})$.

Theorem 4. *Let* $(N, v, \mathcal{S})$ *be a game with coalition structure, let i and j be substitutes, and let* $x \in core(N, v, \mathcal{S})$*. If i and j belong to different members of* $\mathcal{S}$*, then* $x_i = x_j$.

Proof. Let $(i, j) \in N^2$ be substitutes, $\mathcal{C} \in \mathcal{S}$ such that $i \in \mathcal{C}$ and $j \notin \mathcal{C}$. Let $x \in core(N, v, \mathcal{S})$. Since i and j are substitutes, we have $v((\mathcal{C} \setminus \{i\}) \cup \{j\}) = v((\mathcal{C} \setminus \{i\}) \cup \{i\}) = v(\mathcal{C})$. Since $x \in core(N, v, \mathcal{S})$, we have $\forall \mathcal{C} \subseteq N$, $x(\mathcal{C}) \geq v(\mathcal{C})$, we apply this to $(\mathcal{C} \setminus \{i\}) \cup \{j\}$:
$0 \geq v((\mathcal{C} \setminus \{i\}) \cup \{j\}) - x((\mathcal{C} \setminus \{i\}) \cup \{j\}) = v(\mathcal{C}) - x(\mathcal{C}) + x_i - x_j$, since $\mathcal{C} \in \mathcal{S}$ and $x \in core(N, v, \mathcal{S})$, $x(\mathcal{C}) = v(\mathcal{C})$ we have $x_j \geq x_i$. Since i and j play symmetric roles, we have $x_i = x_j$. ✔ □

2.4 The nucleolus

The nucleolus is based on the notion of excess and has been introduced by Schmeidler [79]. The excess measures the amount of "complaints" of a coalition for a payoff distribution. We already mentionned the excess and gave a definition of the core using the excess. We now recall the definition.

Definition 2.13 (Excess)**.** *Let* (N, v) *be a TU game,* $\mathcal{C} \subseteq N$ *be a coalition, and x be a payoff distribution over N. The* excess $e(\mathcal{C}, x)$ *of coalition* $\mathcal{C}$ *at x is the quantity* $e(\mathcal{C}, x) = v(\mathcal{C}) - x(\mathcal{C})$.

The following example in Table 3 is a motivating example for the nucleolus.

A payoff distribution is in the nucleolus when it yields the "least problematic" sequence of complaints according to the lexicographical ordering, i.e., when no other payoff distribution is better. The first entry of $e(x)^{\blacktriangleright}$ is the maximum excess: the agents involved in the corresponding coalition have the largest incentive to leave their current coalition and form a new one. Put another way, the agents involved in that coalition have the most valid complaint. The nucleolus of the game is a set of payoff distributions such that the corresponding vector of excess $e(x)^{\blacktriangleright}$ is minimal. The nucleolus tries to minimise the possible complaints (or minimise the incentives to create a new coalition) over all possible payoff distributions.

$$N = \{1,2,3\},$$
$$v(\{i\}) = 0 \text{ for } i \in \{1,2,3\}$$
$$v(\{1,2\}) = 5, v(\{1,3\}) = 6, v(\{2,3\}) = 6$$
$$v(N) = 8$$

Let us consider two payoff vectors $x = \langle 3,3,2\rangle$ and $y = \langle 2,3,3\rangle$.

$x = \langle 3,3,2\rangle$

coalition $\mathcal{C}$	$e(\mathcal{C},x)$
$\{1\}$	-3
$\{2\}$	-3
$\{3\}$	-2
$\{1,2\}$	-1
$\{1,3\}$	1
$\{2,3\}$	1
$\{1,2,3\}$	0

$y = \langle 2,3,3\rangle$

coalition $\mathcal{C}$	$e(\mathcal{C},y)$
$\{1\}$	-2
$\{2\}$	-3
$\{3\}$	-3
$\{1,2\}$	0
$\{1,3\}$	1
$\{2,3\}$	0
$\{1,2,3\}$	0

Which payoff should we prefer? x or y? To compare two vectors of complaints, we can use the lexicographical order. Let $\geq_{lex}$ denote the *lexicographical* ordering of $\mathbb{R}^m$, i.e., $\forall (x,y) \in \mathbb{R}^m$, $x \geq_{lex} y$ iff

$$\begin{cases} \text{x=y or} \\ \exists t \text{ s. t. } 1 \leq t \leq m \text{ s. t. } \forall i \text{ s. t. } 1 \leq i \leq t\ x_i = y_i \text{ and } x_t > y_t \end{cases}$$

For example, using the lexicographical ordering, we have $\langle 1,1,0,-1,-2,-3,-3\rangle \geq_{lex} \langle 1,0,0,0,-2,-3,-3\rangle$.
Let l be a sequence of m reals. We denote by $l^{\blacktriangleright}$ the *reordering* of l in *decreasing* order.
In the example, $e(x) = \langle -3,-3,-2,-1,1,1,0\rangle$ and then $e(x)^{\blacktriangleright} = \langle 1,1,0,-1,-2,-3,-3\rangle$.
Using the lexicographical ordering, we are now ready to compare the payoff distributions x and y and we note that y is better than x since $e(x)^{\blacktriangleright} \leq_{lex} e(y)^{\blacktriangleright}$: there is a smaller amount of complaints in y than in x given the lexicographical ordering.

Table 3: A motivating example for the nucleolus

Definition 2.14. *Let $\mathcal{I}mp$ be the set of all imputations. The* nucleolus $Nu(N,v)$ *is the set*
$$Nu(N,v) = \left\{x \in \mathcal{I}mp \mid \forall y \in \mathcal{I}mp\ \ e(y)^{\blacktriangleright} \geq_{lex} e(x)^{\blacktriangleright}\right\}.$$

We now provide some properties of the nucleolus.

Theorem 5. *Let (N,v) be a TU game and $\mathcal{I}mp$ is the set of imputations. If $\mathcal{I}mp \neq \emptyset$, then the nucleolus $Nu(N,v)$ is* non-empty.

Theorem 6. *The nucleolus has* at most one *element.*

Theorem 7. *Let (N,v) be a TU game with a non-empty core. Then $Nu(N,v) \subseteq core(N,v)$*

The nucleolus is guaranteed to be non-empty and it is unique. These are two important property in favour of the nucleolus. Moreover, when the core is non-empty, the nucleolus is in the core.

One drawback, however, is that the nucleolus is difficult to compute. It can be computed using a sequence of linear programs of decreasing dimensions. The size of each of these groups is, however, exponential. In some special cases, the nucleolus can be computed in polynomial time [48, 29], but in the general case, computing the nucleolus is not guaranteed to be polynomial. Only a few papers in the multiagent systems community have used the nucleolus, e.g., [94].

2.5 The kernel

The kernel was first introduced by Davis and Maschler in [27]: in the kernel, the strength of the players is measured by the maximum excess the agent can obtain by forming a new coalition with different agents. An agent can consider a payoff distribution to be acceptable by comparing its own 'strength' with the 'strength' of other members of its coalition. When both agents have equal strength, they do not have any incentive to leave the coalition. Although its definition is not as intuitive as the core, the kernel exists and is always non-empty.

2.5.1 Definition of the kernel

We recall that the *excess* related to coalition $\mathcal{C}$ for a payoff distribution x is defined as $e(\mathcal{C},x) = v(\mathcal{C}) - x(\mathcal{C})$. For two agents k and l, the *maximum surplus* $s_{k,l}$ of agent k over agent l with respect to x is $\max_{c\subseteq N \mid k\in c,\, l\notin c} e(c,x)$.

This maximum surplus can be used by agent k to show its strength over agent l: assuming it is positive and that the agent can claim all of it, agent k can argue that it will be better off without agent l; hence it should be compensated with more utility for staying in the current coalition. When any two agents in a coalition have the same maximum surplus (except for a special case), the agents are said to be in equilibrium. A payoff distribution is in the kernel when all agents are in equilibrium. The formal definitions follow:

Definition 2.15 (kernel)**.** *Let $(N,v,\mathcal{S})$ be a TU game with coalition structure. The* kernel *is the set of imputations $x \in \mathcal{I}mp(\mathcal{S})$ such that for every coalition $\mathcal{C} \in \mathcal{S}$, if $(k,l) \in \mathcal{C}^2$, $k \neq l$, then we have either $s_{kl}(x) \geq s_{lk}(x)$ or $x_k = v(\{k\})$.*

One property of the kernel is that the nucleolus is included in the kernel. Hence, it follows that when the set of imputations is non-empty, the kernel is also non empty.

Theorem 8. *The nucleolus is included in the kernel*

Theorem 9. *When $\mathcal{I}mp \neq \emptyset$, then the kernel is non-empty.*

An approximation of the kernel is the ϵ-kernel, where the equality $s_{k,l} = s_{l,k}$ above is replaced by $|s_{k,l} - s_{l,k}| \leq \epsilon$. Notice that this definition requires considering all pairs of agents in the coalition. Unlike the case of the core, the kernel is always non-empty [61]. One property of the kernel is that agents with the same maximum surplus, i.e., symmetric agents, will receive equal payoff. For ensuring fairness, this property is important.

2.5.2 Computational Issues

One method for computing the kernel is the Stearns method [87]. The idea is to build a sequence of side-payments between agents to decrease the difference of surpluses between the agents. At each step of the sequence, the agents with the largest maximum surplus difference exchange utility so as to decrease their surplus: the agent with smaller surplus makes a payment to an agent with higher surplus so as to decrease their surplus difference. After each side-payment, the maximum surplus over all agents decreases. In the limit, the process converges to an element in the kernel. Computing an element in the kernel may require an infinite number of steps as the side payments can become arbitrarily small, and the use of the ϵ-kernel can alleviate this issue. A criteria to terminate Stearns method is proposed in [83], and we present the corresponding algorithm in Algorithm 1.

Algorithm 1: Transfer scheme to converge to a ϵ-kernel-stable payoff distribution for the CS S

```
compute-ε-kernel(ε, S)
repeat
    for each coalition C ∈ S do
        for each member i ∈ C do
            for each member j ∈ C, j ≠ i, do    // compute the surplus for two members of a
            coalition in S
                s_ij ← max_{R⊆N|(i∈R, j∉R)} v(R) − x(R)
    δ ← max_{(i,j)∈N²} |s_ij − s_ji|;
    (i*, j*) ← argmax_{(i,j)∈N²} s_ij − s_ji;
    if (x_j* − v({j}) < δ/2) then                // payment should be individually rational
        d ← x_j* − v({j*});
    else
        d ← δ/2;
    x_i* ← x_i* + d;
    x_j* ← x_j* − d;
until δ/v(S) ≤ ε ;
```

Computing a kernel distribution is of exponential complexity. In Algorithm 1, computing the surpluses is expensive, as we need to search through all coalitions that contains a particular agent and does not contain another agent. Note that when a side-payment is performed, it is necessary to recompute the maximum surpluses. The derivation of the complexity of the Stearns method to compute a payoff in the ϵ-kernel can be found in [45, 83], and the complexity for one side-payment is $O(n \cdot 2^n)$. Of course, the number of side-payments depends on the precision ϵ and on the initial payoff distribution. They derive an upper bound for the number of iterations: converging to an element of the ϵ-kernel requires $n \log_2(\frac{\delta_0}{\epsilon \cdot v(S)})$, where δ_0 is the maximum surplus difference in the initial payoff distribution. To derive a polynomial algorithm, the number of coalitions must be bounded. The solution used in [45, 83] is to only consider coalitions whose size is bounded in the interval K_1, K_2. The complexity of the truncated algorithm is $O(n^2 \cdot n_{coalitions})$ where $n_{coalitions}$ is the number of coalitions with a size between K_1 and K_2, which is a polynomial of order K_2.

2.6 Shapley Value

The Shapley value is designed to provide a fair payoff distribution in a coalition [81]. First we present a set of axioms that defines the Shapley value. We then present a different interpretation of the Shapley value which is based on marginal surplus. Finally, we present computational issues and a value that is cheaper to compute which derived from the Shapley value.

2.6.1 An Axiomatic Characterisation

Another approach to define a stability concept is the axiomatic method. We consider a *value function* ϕ which assigns an efficient allocation x to a TU game (N, v) (i.e., an allocation that satisfies $\phi(N) = v(N)$). Let us consider some axioms that may be desirable for such a value function. The first axiom uses the definition of a dummy agent: an agent i is a *dummy* when $v(\mathcal{C} \cup i) - v(\mathcal{C}) = v(i)$ for all $\mathcal{C} \subseteq N$ such that $i \notin \mathcal{C}$.

DUM ("Dummy actions") : if agent i is a dummy then $x_i = v(\{i\})$. In other words, if the presence of agent i does not improve the worth of a coalition by more than $v(\{i\})$, the agent does not bring anything to the coalition, and then, should obtain only $v(\{i\})$.

SYM ("Symmetry") : When two agents generate the same marginal contributions, they should be rewarded equally: for $i \neq j$ and $\forall \mathcal{C} \subseteq N$ such that $i \notin \mathcal{C}$ and $j \notin S$, if $v(\mathcal{C} \cup \{i\}) = v(\mathcal{C} \cup \{j\})$, then $x_i = x_j$.

ADD ("Additivity") : For any two TU games (N, v) and (N, w) and corresponding payoff profiles $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^n$, the payoff profile should be $x + y$ for the TU game $(N, v + w)$.

Theorem 10. *The Shapley value is the unique value that is budget-balanced and that satisfies axioms 1, 2, and 3.*

The theorem states that these three axioms uniquely define a value function, that is called the *Shapley value*. A proof of this theorem can be found in [61]. To prove this results, one needs to show the existence of a value function that satisfies the three axioms, and then prove the unicity of the value function. In addition, note that the axioms are independent. Finally, one can also show that if one of the three axioms is dropped, it is possible to find multiple value functions satisfying the other two axioms.

The axioms SYM and DUM are clearly desirable. The last axion, ADD, is harder to motivate in some cases. If the valuation function of a TU game is interpreted as an expected payoff, then ADD is desirable. Also, if we consider cost-sharing games and one TU game corresponds to sharing the cost of one service, then ADD is desirable as the cost for a joint-service should be the sum of the cost of the separate services. However, if we do not make any assumptions about the games (N, v) and (N, w), the axiom implies that there is no interaction between the two games. In addition, the game $(N, v + w)$ may induce a behavior that may be unrelated to the behavior induced by either (N, v) or (N, w). Other axiomatisations that do not use the ADD axiom have been proposed by [95] and [58].

2.6.2 Ordinal Marginal Contribution

Another interpretation of the Shapley value is based on the notion of ordered marginal contribution. The marginal contribution of an agent i to a coalition $\mathcal{C} \subseteq N$ is $mc_i(\mathcal{C}) = v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})$. Let us consider that a coalition $\mathcal{C}$ is built incrementally with one agent at a time entering the coalition. Also consider that the payoff of each agent i is its marginal contribution. For example, $\langle mc_1(\emptyset), mc_2(\{1\}), mc_3(\{1,2\})\rangle$ is an efficient payoff distribution for a game $\{1,2,3\}, v)$[1]. In this case, the value of each agent depends on the order in which the agents enter the coalition, which may not be fair. For example, consider agents that form a coalition to take advantages of price reduction when buying large quantities of a product. Agents that start the coalition may have to spend large setup cost, and agents that come later benefits from the already large number of agents in the coalition. To alleviate this issue, the Shapley value averages each agents' payoff over all possible orderings: the value of agent i in coalition $\mathcal{C}$ is the average marginal value over all possible orders in which the agents may join the coalition.

Let σ represent a joining order of the grand coalition N: σ can also be viewed as a permutation of $\langle 1, \ldots, n\rangle$. We write $mc(\sigma)$ the payoff vector where agent i obtains $mc_i(\{\sigma(j) \mid j < i\})$. The payoff vector $mc(\sigma)$ is called

[1]the proof of Theorem 1 shows that such a payoff distribution is in the core of a convex game

the *marginal vector*. Let us denote the set of all permutations of the sequence $\langle 1, \ldots, n\rangle$ as $\Sigma(N)$. The Shapley values can then be defined as

$$Sh(N, v) = \frac{\sum_{\sigma \in \Sigma(N)} mc(\sigma)}{n!}.$$

$$N = \{1,2,3\} \quad \begin{array}{ccc} v(\{1\}) = 0 & v(\{2\}) = 0 & v(\{3\}) = 0 \\ v(\{1,2\}) = 90 & v(\{1,3\}) = 80 & v(\{2,3\}) = 70 \\ & v(\{1,2,3\}) = 120 & \end{array}$$

	1	2	3
$1 \leftarrow 2 \leftarrow 3$	0	90	30
$1 \leftarrow 3 \leftarrow 2$	0	40	80
$2 \leftarrow 1 \leftarrow 3$	90	0	30
$2 \leftarrow 3 \leftarrow 1$	50	0	70
$3 \leftarrow 1 \leftarrow 2$	80	40	0
$3 \leftarrow 2 \leftarrow 1$	50	70	0
total	270	240	210
Shapley value ϕ	45	40	35

Let $y = \langle 50, 40, 30\rangle$

$\mathcal{C}$	$e(\mathcal{C}, \phi)$	$e(\mathcal{C}, y)$
$\{1\}$	-45	0
$\{2\}$	-40	0
$\{3\}$	-35	0
$\{1,2\}$	5	0
$\{1,3\}$	0	0
$\{2,3\}$	-5	0
$\{1,2,3\}$	0	0

This example shows that the Shapley value may not be in the core, and may not be the nucleolus.

Table 4: Example of a computation of Shapley value

We provide an example in Table 4 in which we list all the orders in which the agents can enter the grand coalition. The sum is over all joining orders, which may contains a very large number of terms. However, when computing the Shapley value for one agent, one can avoid some redundancy by summing over all coalitions and noticing that:

- There are $|\mathcal{C}|!$ permutations in which all members of $\mathcal{C}$ precede i.
- There are $|N \setminus (\mathcal{C} \cup \{i\})|!$ permutations in which the remaining members succede i, i.e. $(n - |\mathcal{C}| - 1)!$.

These observations allow us to rewrite the Shapley value as:

$$Sh_i(N, v) = \sum_{\mathcal{C} \subseteq N \setminus \{i\}} \frac{|\mathcal{C}|!(n - |\mathcal{C}| - 1)!}{n!} \left(v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})\right).$$

Note that the example from Table 4 also demonstrates that in general the Shapley value is not in the core or in the nucleolus. At noted before, the Shapley value always exists and is unique. When the valuation function is superadditive, the Shapley value is individually rational, i.e., it is an imputation. When the core is non-empty, the Shapley value may not be in the core. However, when the valuation function is convex, the Shapley value is also group rational, hence, it is in the core.

2.6.3 Computational Issues

The nature of the Shapley value is combinatorial, as all possible orderings to form a coalition needs to be considered. This computational complexity can sometimes be an advantage as agents cannot benefit from manipulation. For example, it is $\mathcal{NP}$-complete to determine whether an agent can benefit from false names [93]. Nevertheless, the structure of the characteristic function can be exploited to rapidly compute the Shapley value. Conitzer and Sandholm consider characteristic functions that can be decomposed over multiple issues that interest only a subset of the agents [22]. When few agents are concerned by each issue, and the number of issues is small, fast computation of the Shapley value is possible. In [40], the characteristic function is represented by a set of rules and the computation of the Shapley value can be performed in time linear in the number of rules. We now provide more details about these approaches.

Conitzer and Sandholm analyse the case where the agents are concerned with multiple independent issues that a coalition can address. For example, performing a task may require multiple abilities, and a coalition may gather agents that work on the same task but with limited or no interactions between them. A characteristic function v can be decomposed over T issues when it is of the form $v(\mathcal{C}) = \sum_{t=1}^{T} v_t(\mathcal{C})$). In that case, the Shapley value $Sh_i(\mathcal{C})$ for agent i for the characteristic function v is the sum of the Shapley value $Sh_i^t(\mathcal{C})$ for agent i over the t different issues: $Sh_i(\mathcal{C}) = \sum_{t=1}^{T} Sh_i^t(\mathcal{C}, i)$. If only a small number of agents is concerned about an issue, computing the

Shapley value for the particular issue can be cheap. For an issue t, the characteristic function v_t concerns only the agents in $\mathcal{I}_t$ when $\forall \mathcal{C}_1 \subseteq N, \mathcal{C}_2 \subseteq N$ such that $\mathcal{I}_t \cap \mathcal{C}_1 = \mathcal{I}_t \cap \mathcal{C}_2 \Rightarrow v_t(\mathcal{C}_1) = v_t(\mathcal{C}_2)$. When the characteristic function v is decomposed over T issues and when $|\mathcal{I}_t|$ agents are concerned about each issue $t \in [1...T]$, computing the Shapley value takes $O(\sum_{t=1}^{T} 2^{|\mathcal{I}_t|})$.

The characteristic function can also be represented by a set of "rules" [40]. Each "rule" associates a pattern and a value. The pattern is a conjunction of agents. A coalition matches the pattern if it is a super-set of the pattern. The value of a coalition is the sum over the values of all the rules that apply to the coalition. The pattern can be extended to positive or negative literals for representational efficiency: a positive literal represents the presence of an agent in a coalition, whereas a negative literal represents the absence of an agent in the coalition. This representation can be much more concise than the traditional representation of the characteristic function for certain games. Under this representation, the algorithm for computing the Shapley value runs in linear time of the input.

2.6.4 Bilateral Shapley Value

In order to reduce the combinatorial complexity of the computation of the Shapley value, Ketchpel introduces the Bilateral Shapley Value (BSV) [42]. The idea is to consider the formation of a coalition as a succession of merging between two coalitions. Two disjoint coalitions $\mathcal{C}_1$ and $\mathcal{C}_2$ with $\mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$, may merge when $v(\mathcal{C}_1 \cup \mathcal{C}_2) \geq v(\mathcal{C}_1) + v(\mathcal{C}_2)$. When they merge, the two coalitions, called founders of the new coalition $\mathcal{C}_1 \cup \mathcal{C}_2$, share the marginal utility as follows: $BSV(\mathcal{C}_1) = \frac{1}{2}v(\mathcal{C}_1) + \frac{1}{2}(v(\mathcal{C}_1 \cup \mathcal{C}_2) - v(\mathcal{C}_2))$ and $BSV(\mathcal{C}_2) = \frac{1}{2}v(\mathcal{C}_2) + \frac{1}{2}(v(\mathcal{C}_1 \cup \mathcal{C}_2) - v(\mathcal{C}_1))$. This is the expression of the Shapley value in the case of an environment with two agents. In $\mathcal{C}_1 \cup \mathcal{C}_2$, each of the founders gets half of its 'local' contribution, and half of the marginal utility of the other founder. Given this distribution of the marginal utility, it is rational for $\mathcal{C}_1$ and $\mathcal{C}_2$ to merge if $\forall i \in \{1, 2\}$, $v(\mathcal{C}_i) \leq BSV(C_i)$. Note that symmetric founders get equal payoff, i.e., for $\mathcal{C}_1$, $\mathcal{C}_2$, $\mathcal{C}$ such that $\mathcal{C}_1 \cap \mathcal{C}_2 = \mathcal{C}_1 \cap \mathcal{C} = \mathcal{C}_2 \cap \mathcal{C} = \emptyset$, $v(\mathcal{C} \cup \mathcal{C}_1) = v(\mathcal{C} \cup \mathcal{C}_2) \Rightarrow BSV(\mathcal{C} \cup \mathcal{C}_1) = BSV(\mathcal{C} \cup \mathcal{C}_2)$. Given a sequence of successive merges from the states where each agent is in a singleton coalition, we can use a backward induction to compute a stable payoff distribution [44]. Though the computation of the Shapley value requires looking at all of the permutations, the value obtained by using backtracking and the BSV only focuses on a particular set of permutations, but the computation is significantly cheaper.

2.7 Voting Games

The formation of coalitions is usual in parliaments or assemblies. It is therefore interesting to consider a particular class of coalitional games that models voting in an assembly. For example, we can represent an election between two candidates as a voting game where the winning coalitions are the coalitions of size at least equal to the half the number of voters. The formal definition follows:

Definition 2.16 (voting game)**.** *A game (N, v) is a* voting game *when*

- *the valuation function takes only two values: 1 for the winning coalitions, 0 otherwise.*
- *v satisfies* unanimity*: $v(N) = 1$*
- *v satisfies* monotonicity*: $S \subseteq T \subseteq N \Rightarrow v(S) \leq v(T)$.*

Unanimity and monotonicity are natural assumptions in most cases. Unanimity reflects the fact that all agents agree; hence, the coalition should be winning. Monotonicity tells that the addition of agents in the coalition cannot turn a winning coalition into a losing one, which is reasonable for voting: more supporters should not harm the coalition. A first way to represent a voting game is by listing all winning coalitions. Using the monotonicity property, a more succint representation is to list only the *minimal winning coalitions*.

Definition 2.17 (Minimal winning coalition)**.** *A coalition $\mathcal{C} \subseteq N$ is a minimal winning coalition iff $v(\mathcal{C}) = 1$ and $\forall i \in \mathcal{C}$ $v(\mathcal{C} \setminus \{i\}) = 0$.*

We can now see how we formalize some common terms in voting: dictatorship, veto player and blocking coalition.

Definition 2.18 (Dictator)**.** *Let (N, v) be a simple game. A player $i \in N$ is a* dictator *iff $\{i\}$ is a winning coalition.*

Note that with the requirements of simple games, it is possible to have more than one dictator!

Definition 2.19 (Veto Player)**.** *Let (N, v) be a simple game. A player $i \in N$ is a* veto *player if $N \setminus \{i\}$ is a losing coalition. Alternatively, i is a* veto *player iff for all winning coalition $\mathcal{C}$, $i \in \mathcal{C}$.*

It also follows that a veto player is member of every minimal winning coalitions.

Definition 2.20 (blocking coalition)**.** *A coalition $\mathcal{C} \subseteq N$ is a* blocking coalition *iff $\mathcal{C}$ is a losing coalition and $\forall S \subseteq N \setminus \mathcal{C}$, $S \setminus \mathcal{C}$ is a losing coalition.*

A variant of a voting game is a weighted voting game where each agent has a weight and a coalition needs to achieve a threshold or quota to be winning. This is a much more compact representation as we only use to define a vector of weights and a threshold. However, this is not a complete representations are some voting games cannot be represented as a weighted voting game. The formal definition follows.

Definition 2.21 (weighted voting game)**.** *A game (N, v, q, w) is a* weighted voting game *when*

- *$w = (w_1, w_2 \ldots, w_n)$ is a vector of weights, one for each voter*
- *A coalition $\mathcal{C}$ is winning (i.e., $(v(\mathcal{C}) = 1)$ iff $\sum_{i \in \mathcal{C}} w_i \geq q$, it is losing otherwise (i.e., $(v(\mathcal{C}) = 0)$*
- *v satisfies monotonicity: $S \subseteq T \subseteq N \Rightarrow v(S) \leq v(T)$.*

We will note a weighted voting game $(N, w_{i \in N}, q)$ as $[q;\ w_1, \ldots, w_n]$. Note that the weight may not represent the voting power of the player. Let us consider the following weighted voting games:

- $[10;\ 7, 4, 3, 3, 1]$: The set of minimal winning coalitions is $\{\{1, 2\}\{1, 3\}\{1, 4\}\{2, 3, 4\}\}$. Player 5, although it has some weight, is a dummy. Player 2 has a higher weight than player 3 and 4, but it is clear that player 2, 3 and 4 have the same influence.
- $[51;\ 49, 49, 2]$: The set of winning coalition is $\{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$. It seems that the players have symmetric roles, but it is not reflected in their weights.

The European Union uses a combination of weighted voting games (a decision is accepted when it is supported by 55% of Member States, including at least fifteen of them, representing at the same time at least 65% of the Union's population).

The examples raise the subject of measuring the voting power of the agents in a voting game. Multiple indices have been proposed to answer these questions, and we now present few of them. One central notion is the notion of *pivotal player*: we say that a voter i is pivotal for a coalition $\mathcal{C}$ when it turns it from a losing to a wining coalition, i.e., $v(\mathcal{C}) = 0$ and $v(\mathcal{C} \cup \{i\}) = 1$. Let w be the number of winning coalitions. For a voter i, let η_i be the number of coalitions for which i is pivotal, i.e., $\eta_i = \sum_{S \subseteq N \setminus \{i\}} v(S \cup \{i\}) - v(S)$.

Shapley-Shubik index: it is the Shapley value of the voting game, its interpretation in this context is the percentage of the permutations of all players in which i is pivotal.

$$I_{SS}(N, v, i) = \sum_{\mathcal{C} \subseteq N \setminus \{i\}} \frac{|\mathcal{C}|!(n - |C| - 1)!}{n!} \left(v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})\right).$$

"For each permutation, the pivotal player gets one more point."

One issue is that the voters do not trade the value of the coalition, though the decision that the voters vote about is likely to affect the entire population.

Banzhaff index: For each coalition, we determine which agent is a swing agent (more than one agent may be pivotal). The *raw Banzhaff index* of a player i is

$$\beta_i = \frac{\sum_{\mathcal{C} \subseteq N \setminus \{i\}} v(\mathcal{C} \cup \{i\}) - v(\mathcal{C})}{2^{n-1}}.$$

For a simple game (N, v), $v(N) = 1$ and $v(\emptyset) = 0$, at least one player i has a power index $\beta_i \neq 0$. Hence, $B = \sum_{j \in N} \beta_j > 0$. The *normalized Banzhaff index* of player i for a simple game (N, v) is defined as

$$I_B(N, v, i) = \frac{\beta_i}{B}.$$

We provide in Table 5 an example of computation of the Shapley-Schubik and Banzhaff indices. This example shows that both indices may be different.

The computational complexity of voting and weighted voting games have been studied in [30, 33]. For example, the problem of determining whether the core is empty is polynomial. When the core is non-empty, the problem of computing the nucleolus is also polynomial, otherwise, it is an $\mathcal{NP}$-hard problem. The problem of choosing the weights so that they corresponds to a given power index has also been tackled in [28].

$\{1,\underline{2},3,4\}$ $\{3,1,\underline{2},4\}$
$\{1,\underline{2},4,3\}$ $\{3,1,\underline{4},2\}$
$\{1,3,\underline{2},4\}$ $\{3,2,\underline{1},4\}$
$\{1,3,\underline{4},2\}$ $\{3,2,4,\underline{1}\}$
$\{1,4,\underline{2},3\}$ $\{3,4,\underline{1},2\}$
$\{1,4,\underline{3},2\}$ $\{3,4,2,\underline{1}\}$
$\{2,\underline{1},3,4\}$ $\{4,1,\underline{2},3\}$
$\{2,\underline{1},4,3\}$ $\{4,1,\underline{3},2\}$
$\{2,3,\underline{1},4\}$ $\{4,2,\underline{1},3\}$
$\{2,3,4,\underline{1}\}$ $\{4,2,3,\underline{1}\}$
$\{2,4,\underline{1},3\}$ $\{4,3,\underline{1},2\}$
$\{2,4,3,\underline{1}\}$ $\{4,3,2,\underline{1}\}$
In red and underlined, the pivotal agent

	1	2	3	4
Sh	$\frac{7}{12}$	$\frac{1}{4}$	$\frac{1}{12}$	$\frac{1}{12}$

winning coalitions:
$\{\underline{1},\underline{2}\}$
$\{\underline{1},\underline{2},3\}$
$\{\underline{1},\underline{2},4\}$
$\{\underline{1},\underline{3},\underline{4}\}$
$\{\underline{1},2,3,4\}$
In red and underlined, the pivotal agents

	1	2	3	4
β	$\frac{5}{8}$	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{1}{8}$
$I_B(N,v,i)$	$\frac{1}{2}$	$\frac{3}{10}$	$\frac{1}{10}$	$\frac{1}{10}$

Table 5: Shapley-Schubik and the Banzhaff indices for the weighted voting game $[7;\ 4,3,2,1]$.

2.8 Game with externalities

A traditional assumption in the literature of coalition formation is to consider that the value of a coalition does not depend on non-members' actions. In general, this may not be true: some externalities (positive or negative) can create a dependency between the value of a coalition and the actions of non-members. Sandholm and Lesser attribute these externalities to the presence of shared resources (if a coalition uses some resource, they will not be available to other coalitions), or when there are conflicting goals: non-members can move the world farther from a coalition's goal state [75]. Ray and Vohra state in [74] that a "recipe for generating characteristic functions is a minimax argument": the value of a coalition $\mathcal{C}$ is the value $\mathcal{C}$ gets when the non-members respond optimally so as to minimise the payoff of $\mathcal{C}$. This formulation acknowledges that the presence of other coalitions in the population may affect the payoff of the coalition $\mathcal{C}$. As in [37, 74], we can study the interactions between different coalitions in the population: decisions about joining forces or splitting a coalition can depend on the way the competitors are organised. For example, when different companies are competing for the same market niche, a small company might survive against a competition of multiple similar individual small companies. However, if some of these small companies form a viable coalition, the competition significantly changes: the other small companies may now decide to form another coalition to be able to successfully compete against the existing coalition. Another such example is a bargaining situation where agents need to negotiate over the same issues: when agents form a coalition, they can have a better bargaining position, as they have more leverage, and because the other party needs to convince all the members of the coalition. If the other parties also form coalition, the bargaining power of the first coalition may decrease.

We believe that forming coalitions in domains where the value depends on the CS is an important problem which deserves more attention. Because of the dependencies of the value and the CS, it is also a harder problem. Recently the topic has raised interest in AI. Rahwan et al. consider the problem of CS generation in this case (we will present this problem later) [70]. Michalak et al. tackle the problem of representing such games (they use a more compact description, still allowing efficient computation)[57] . Finally, Elkind et al. consider the restriction for weighted voting games [32]. In the following, we introduce two important models: the games in partition function form and the games with valuations.

2.8.1 Games in partition function form

A first model is the partition function form in which the value of a coalition depends on the CS it belongs.

Definition 2.22 (Games in partition function form)**.** *A game in partition function form is a game (N,v) where the valuation function $v : 2^N \times \mathscr{S}_N \to \mathbb{R}$.*

The definition of such a function requires an very large space. As shown in our example with the different firms, the effect of a merge of two coalitions on the remaining coalition in the CS is key. We can consider two important families: one for which the effect of the merge is positive on the remaining coalition (positive spillover), and one for which the effect is negative.

Definition 2.23 (Positive and negative spillovers)**.** *A partition function v exhibits* positive *(resp.* negative*)* spillovers *if for any CS $\mathcal{S}$ and any two coalitions C_1 and C_2 in $\mathcal{S}$, $v(\mathcal{C},\mathcal{S}\setminus\{\mathcal{C}_1,\mathcal{C}_2\}\cup\{\mathcal{C}_1\cup\mathcal{C}_2\}) \geq v(C,\mathcal{S})$ (resp. $v(\mathcal{C},\mathcal{S}\setminus\{\mathcal{C}_1,\mathcal{C}_2\}\cup\{\mathcal{C}_1\cup\mathcal{C}_2\}) \leq v(C,\mathcal{S})$) for all coalitions $\mathcal{C}\neq\mathcal{C}_1,\mathcal{C}_2$ in $\mathcal{S}$.*

2.8.2 Games with valuations

Fixed rules of division appear naturally in many economic situations and in theoretical studies based on a two-stage procedure where first the coalitions are formed and only then the payoff distribution problem is solved. Hence, a second model considers games which are directly defined in terms of a payoff distribution that depends on the CS.

Definition 2.24 (Games in Valuation)**.** *A* game in valuation (N,v) *is a game for which the valuation function is a mapping $v : \mathscr{S}_N \to \mathbb{R}^n$ which associates to each coalition structure a payoff distribution in $\mathbb{R}^n$.*

We can define the notion of positive and negative spillovers for these games as follows:

Definition 2.25 (Positive and negative spillovers)**.** *A valuation v exhibits* positive *(resp.* negative*)* spillovers *if for any CS $\mathcal{S}$ and any two coalitions $\mathcal{C}_1$ and $\mathcal{C}_2$ in $\mathcal{S}$, $v_i(\mathcal{S}\setminus\{\mathcal{C}_1,\mathcal{C}_2\}\cup\{\mathcal{C}_1\cup\mathcal{C}_2\}) \geq v_i(\mathcal{S})$ (resp. $v_i(\mathcal{S}\setminus\{\mathcal{C}_1,\mathcal{C}_2\}\cup\{\mathcal{C}_1\cup\mathcal{C}_2\}) \leq v_i(\mathcal{S})$) for all players $i\notin\mathcal{C}_1\cup\mathcal{C}_2$.*

3 Non-Transferable Utility Games (NTU games)

The underlying assumption behind a TU game is that agents have a common scale to measure the worth of a coalition. Such a scale may not exist in every situation, which leads to the study of games where the utility is non-transferable (NTU games). We start by introducing a particular type of NTU games called *Hedonic games*.

In these games, agents have preferences over coalitions, i.e., agents value the company of the other members of the coalition, hence the name. Let N be a set of agents and $\mathcal{N}_i$ be the set of coalitions that contain agent i, i.e., $\mathcal{N}_i = \{\mathcal{C}\cup\{i\} \mid \mathcal{C}\subseteq N\setminus\{i\}\}$. For a CS $\mathcal{S}$, we will note $\mathcal{S}(i)$ the coalition in $\mathcal{S}$ containing agent i.

Definition 3.1 (Hedonic games)**.** *An* Hedonic game *is a tuple $(N,(\succeq_i)_{i\in N})$ where*

- *N is the set of agents*
- *$\succeq_i \subseteq 2^{\mathcal{N}_i}\times 2^{\mathcal{N}_i}$ is a complete, reflexive and transitive preference relation for agent i, with the interpretation that if $S\succeq_i T$, agent i prefers coalition T at most as much as coalition S.*

We now give the definition of stability concepts adapted from [16].

Core stability: A CS $\mathcal{S}$ is $\boldsymbol{core\text{-}stable}$ iff $\nexists\mathcal{C}\subseteq N \mid \forall i\in\mathcal{C},\mathcal{C}\succ_i\mathcal{S}(i)$.

Nash stability: A CS $\mathcal{S}$ is $\boldsymbol{Nash\text{-}stable}$ iff $(\forall i\in N)$ $(\forall\mathcal{C}\in\mathcal{S}\cup\{\emptyset\})$ $\mathcal{S}(i)\succsim_i\mathcal{C}\cup\{i\}$. No player would like to join any other coalition in $\mathcal{S}$ assuming the other coalitions did not change.

Individual stability A CS $\mathcal{S}$ is $\boldsymbol{individually\ stable}$ iff $(\nexists i\in N)$ $(\nexists\mathcal{C}\in\mathcal{S}\cup\{\emptyset\}) \mid (\mathcal{C}\cup\{i\}\succ_i\mathcal{S}(i))$ *and* $(\forall j\in\mathcal{C},\ \mathcal{C}\cup\{i\}\succsim_j\mathcal{C})$. No player can move to another coalition that it prefers without making some members of that coalition unhappy.

Contractually individual stability: A CS $\mathcal{S}$ is $\boldsymbol{contractually\ individually\ stable}$ iff $(\nexists i\in N)$ $(\nexists\mathcal{C}\in\mathcal{S}\cup\{\emptyset\}) \mid (\mathcal{C}\cup\{i\}\succ_i\mathcal{S}(i))$ *and* $(\forall j\in\mathcal{C},\ \mathcal{C}\cup\{i\}\succsim_j\mathcal{C})$ *and* $(\forall j\in\mathcal{S}(i)\setminus\{i\},\ \mathcal{S}(i)\setminus\{i\}\succsim_j\mathcal{S}(i))$. No player can move to a coalition it prefers so that the members of the coalition it leaves and it joins are better off.

If a CS is core-stable, no subset of agents has incentive to leave its respective coalition to form a new one. In a Nash-stable CS $\mathcal{S}$, no single agent i has an incentive to either leave its coalition $\mathcal{S}(i)$ to join an existing coalition in $\mathcal{S}$ or to create the singleton coalition $\{i\}$. The two other criteria add a constraint on the members of the coalition joined or left by the agent. For an individually stable CS, there is no agent that can change coalition from $\mathcal{S}(i)$ to $\mathcal{C}\in\mathcal{S}$ yielding better payoff for itself, and the members of $\mathcal{C}$ should not lose utility. The contractually individual stability requires that in addition, the other members of $\mathcal{S}(i)$, the coalition left by the agent i, should not lose utility.

The literature in game theory focuses on finding conditions for the existence of the core. In the AI literature, Elkind and Wooldridge have proposed a succinct representation of Hedonic games [34].

The general definition of an NTU game uses a set of outcomes that can be achieved by the coalitions. The formal definition is the following:

Definition 3.2 (NTU Game)**.** *A non-transferable utility game (NTU Game) $(N,X,V,(\succ_i)_{i\in N})$ is defined by*

- *a set of agents N;*
- *a set of outcomes X;*
- *a function $V : 2^N \rightarrow 2^X$ that describes the outcomes $V(\mathcal{C}) \subseteq X$ that can be brought about by coalition $\mathcal{C}$;*
- *a preference relation $\succ_i$ (transitive and complete) over the set of outcomes for each agent i.*

Intuitively, $V(\mathcal{C})$ is the set of outcomes that $\mathcal{C}$ can bring about by means of its join action. The agents have a preference relation over the outcomes.

Example 1: hedonic games as a special class of NTU games. Let $(N, (\succeq_i^H)_{i \in N})$ be a hedonic game.

- For each coalition $\mathcal{C} \subseteq N$, create a unique outcome $x_\mathcal{C}$.
- For any two outcomes x_S and x_T corresponding to coalitions S and T that contains agent i, We define $\succeq_i$ as follows: $x_S \succeq_i x_T$ iff $S \succeq_i^H T$.
- For each coalition $\mathcal{C} \subseteq N$, we define $V(\mathcal{C})$ as $V(\mathcal{C}) = \{x_C\}$.

Example 2: a TU game can be viewed as an NTU game. Let (N, v) be a TU game.

- We define X to be the set of all allocations, i.e., $X = \mathbb{R}^n$.
- For any two allocations $(x, y) \in X^2$, we define $\succeq_i$ as follows: $x \succeq_i y$ iff $x_i \geq y_i$.
- For each coalition $\mathcal{C} \subseteq N$, we define $V(\mathcal{C})$ as $V(\mathcal{C}) = \{x \in \mathbb{R}^n \mid \sum_{i \in N} x_i \leq v(\mathcal{C})\}$. $V(\mathcal{C})$ lists all the feasible allocation for the coalition $\mathcal{C}$.

First, we can note that the definition of the core can easily be modified in the case of NTU games.

Definition 3.3. $core(V) = \{x \in V(N) \mid \nexists \mathcal{C} \subset N, \nexists y \in V(\mathcal{C}), \forall i \in \mathcal{C}\ y \succ_i x\}$

An outcome $x \in X$ is blocked by a coalition $\mathcal{C}$ when there is another outcome $y \in X$ that is preferred by all the members of $\mathcal{C}$. An outcome is then in the core when it can be achieved by the grand coalition and it is not blocked by any coalition. As is the case for TU game, it is possible that the core of an NTU game is empty.

4 The Cooperative Case: Sharing the Computation of Coalition values and Searching of the Optimal Coalition Structure

In the previous sections, the focus was on individual agents that are concerned with their individual payoff. In this section, we consider TU games (N, v) in which agents are concerned only about the society's payoff: the agents' goal is to maximise utilitarian social welfare. The actual payoff of the agents or the value of its coalition is not of importance in this setting, only the total value generated by the population matters. This is particularly interesting for multiagent systems designed to maximize some objective functions. For example, a rescue operation, it is not important which rovers is helping the most victims, the goal of the rovers is to help as many victims as possible as a whole. In the following, an optimal CS denotes a CS with maximum social welfare.

The space $\mathscr{S}_N$ of all CSs can be represented by a lattice, and an example for a population of four agents is provided in Figure 4. The first level of the lattice consists only of the CS corresponding to the grand coalition $N = \{1, 2, 3, 4\}$, the last level of the lattice contains CS containing singletons only, i.e., coalitions containing a single member. Level i contains all the CSs with exactly i coalitions. The number of CSs at level i is $\mathcal{S}(n, i)$, where $\mathcal{S}$ is the Stirling Number of the Second Kind[2]. The Bell number, $\mathscr{B}(n)$, represents the total number of CSs with n agents, $\mathscr{B}(n) = \sum_{i=0}^{n} \mathcal{S}(n, k)$. This number grows exponentially, as shown in Figure 5, and is $O(n^n)$ and $\omega(n^{\frac{n}{2}})$ [76]. When the number of agents is relatively large, e.g., $n \geq 20$, exhaustive enumeration may not be feasible.

The actual issue is the search of the optimal CS. [76] show that given a TU game (N, v), the finding the optimal CS is an $\mathcal{NP}$-complete problem. In the following, we will consider centralised search where a single agent is performing the search as well as the more interesting case of decentralised search where all agents make the search at the same time on different parts of the search space. Before doing so, we review some work where the valuation function v is not known in advance. In a real application, these values need to be computed; and this may be an issue on its own if the computations are hard, as illustrated by an example by [75] where the computation of a value requires to solve a traveling salesman problem.

[2] $\mathcal{S}(n, m)$ is the number of ways of partitioning a set of n elements into m non-empty sets.

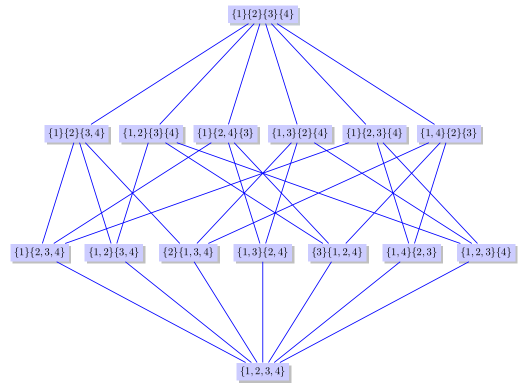

Figure 4: Set of CSs for 4 agents.

4.1 Sharing the computation of the coalition values

Before searching the space of CSs to find an optimal CS, the agents may need to compute the value of each of the coalitions. We are interested in a decentralised algorithm that computes all the coalition values in a minimal amount of time, and that requires minimum communication between the agents.

Shehory and Kraus were the first to propose an algorithm to share the computation of the coalition values [82]. In their algorithm, the agents negotiate which computation is performed by which agent, which is quite demanding. Rahwan and Jennings proposed an algorithm where agents, once they agree on an identification for each agent participating in the computation, know exactly which coalition values to compute. This algorithm, called DCVC [67] outperforms the one by Sheohory and Kraus. The key observation is that in general, it should take longer to compute the value of a large coalition compared to a small coalition (i.e., the computational complexity is likely to increase with the size of the coalition since more agents have to coordinate their activities). Their method improves the balance of the loads by distributing coalitions of the same size to all agents. By knowing the number of agents n participating in the computation an index number (i.e., an integer in the range $\{0..n\}$), the agents determine for each coalition size which coalition values to compute. The algorithm can also be adapted when the agents have different known computational speed so as to complete the computation in a minimum amount of time.

4.2 Searching for the optimal coalition structure

The difficulty of searching for the optimal CS lies in the large search space, as recognised by existing algorithms, and this is even more true in the case where there exists externalities (i.e., when the valuation of a coalition depends on the CS). For TU games with no externalities, some algorithms guarantee finding CSs within a bound from the optimum when an incomplete search is performed. Unfortunately, such guarantees are not possible for games with externalities. We shortly discuss these two cases in the following.

4.2.1 Games with no externalities

[76] proposed a first algorithm that searches through a lattice as presented in Figure 4. Their algorithm guarantees that the CS found, s, is within a bound from the optimal $s^\star$ when a sufficient portion of the lattice has been visited. The bound considered is $\frac{v(s)}{V(s^\star)} \leq \mathcal{K}$. They prove that to ensure a bound, it is necessary to visit a least 2^{n-1} CSs (Theorems 1 and 3 in [76]) which corresponds to the first two levels of the lattice, i.e., the algorithm needs to visit the grand coalition and all the CSs composed of 2 coalitions. The bound improves each time a new level is visited. An empirical study of different strategies for visiting the other levels is presented in [49]. Three different algorithms are empirically tested over characteristic functions with different properties: 1) subadditive, 2) superadditive, 3) picked from a uniform distribution in $[0, 1]$ or in $[0, |\mathcal{S}|]$ (where $|\mathcal{S}|$ is the size of the coalition). The performance of the heuristics differs over the different type of valuation functions, demonstrating the importance of the properties of the characteristic function in the performance of the search algorithm.

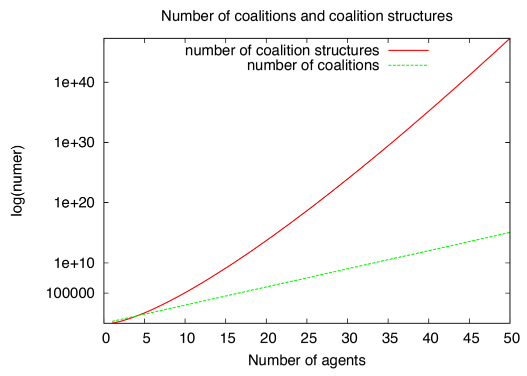

Figure 5: Number of CSs in a population of n agents.

The algorithm by Dang and Jennings [26] improves the one of [76] for low bounds from the optimal. For large bounds, both algorithms visit the first two levels of the lattice. Then, when the algorithm by Sandholm et al. continues by searching each level of the lattice, the algorithm of Dang and Jennings only searches specific subset of each level to decrease the bound faster. This algorithm is anytime, but its complexity is not polynomial.

These algorithms were based on a lattice as the one presented in Figure 4 where a CS in level i contains exactly i coalitions. The best algorithm to date has been developed by Rahwan et al. and uses a different representation called integer-partition (IP) of the search space. It is an anytime algorithm that has been improved over a series of paper: [71, 72, 68, 69, 73]. In this representation the CSs are grouped according to the sizes of the coalitions they contain, which is called a configuration. For example, for a population of four agents, the configuration $\{1, 3\}$ represents CSs that contain a coalition with a singleton and a coalition with three agents. A smart scan of the input allows to search the CSs with two coalitions the grand coalition and the CS containing singletons only. In addition, during the scan, the algorithm computes the average and maximum value for each coalition size. The maximum values can be used to prune the search space. When constructing a configuration, the use of the maximum values of a coalition for each size permits the computation of an upper bound of the value of a CS that follows that configuration, and if the value is not greater than the current best CS, it is not necessary to search through the CSs with that configuration, which prunes the search tree. Then, the algorithm searches the remaining configurations, starting with the most promising ones. During the search of a configuration, a branch and bound technique is used. In addition, during the search, the algorithm is designed so that no CS is evaluated twice. Empirical evaluation shows that the algorithm outperforms any other current approach over different distributions used to generate the values of the coalitions.

4.2.2 Games with externalities

The previous algorithm explicitly uses the fact that the valuation function only depends on the members of the coalition, i.e., has no externalities. When this is not the case, i.e., when the valuation function depends on the CS, it is still possible to use some algorithms, e.g., the one proposed in [49], but the guarantee of being within a bound from the optimal is no longer valid. Sen and Dutta use genetic algorithms techniques [80] to perform the search. The use of such technique only assumes that there exists some underlying patterns in the characteristic function. When such patterns exist, the genetic search makes a much faster improvement in locating higher valued CS compared to the level-by-level search approach. One downside of the genetic algorithm approach is that there is no optimality guarantee. Empirical evaluation, however, shows that the genetic algorithm does not take much longer to find a solution when the value of a coalition does depend on other coalitions.

More recently, Rahwan et al. and Michalak et al. consider the problem for some class of externalities and modify the IP algorithm for the games with externalities [56, 70].

5 Motivating examples from multiagent systems litterature

We now review some domains that were introduced to motivate research about coalition formation in the multiagent systems community. Two domains were exploited: the task allocation domain and electronic marketplace. For both of them, we provide the general application domain and some variants that have been introduced in the literature. These application domains highlight some important issues for artificial agent societies, and we will come back to these issues in Section 6.

5.1 Task Allocation Problem

A task allocation problem can be easily represented by a coalition formation problem: a coalition of agents is in charge of performing a task (or a subset of tasks). A task may require multiple agents to be performed due to the following reasons:

- Complementary expertise may be required to perform a complex task, and many approaches assume that no agent has all the required expertise to perform a complex task on its own [46, 47, 54, 82]. In the general case, a task can be decomposed into subtasks, and the agents are able to perform a subset of all possible subtasks.
- All the agents have the required ability or expertise to perform a task, but they do not have enough resources on their own to perform the task. For example, robots have the ability to move objects in a plant, but multiple robots are required to move a heavy box [5, 82].

In addition, the valuation function of a coalitional game has a simple interpretation: it is the benefit of the group of coalition when the task is performed. The classical stability problem of coalitional games appears since multiple coalitions may be able to perform a complex task, and some coalitions may be better suited to perform a given task. Ideally, an agent should not have incentive to join a different coalition to work on a different set of tasks.

A generic task allocation problem can be described as follows: a coalition of agents forms to perform a complex task and each agent in the coalition plays a role in the completion of the task (they can all have the same or complementary roles). The completion of a task is rewarded by a payoff. The cost associated with the task completion depends on the coalition members. The value of the coalition is the net benefit (payoff minus cost) of completing the task. Hence, the task allocation problem is well-modeled by a coalition formation problem where the value of a coalition depends only on its members. Note that in the case where agents are not self-interested, the population of agents as a whole may try to maximise the total benefit of completing the tasks. In this case, the agents are trying to optimise utilitarian social welfare and search for the optimal CS (see Section 4.2).

Task allocation problems may be even more complex. First, the tasks may be inter-dependent. For example in [82], there is a partial precedence order between the tasks. This assumption is of particular importance in the transportation domain. The existence of task dependency may promote cooperation between the agents as advocated in [6]: the dependence between the tasks may translate into a certain form of dependence between the agents. If agents realise about it, they may reciprocally help each other: agent A may help agent B to perform a task needed for the completion of an important task for agent B, and vice versa.

In the generic model, the cost and benefit depends only on the members of the coalition. In environments where the value of a task depends on its completion time, Kraus et al. suggested that there should be a a cost associated with time it takes to decide on a coalition [46, 47]. They propose a variant of the task allocation problem where at each round, the reward to perform a task is reduced. This forces the agents to decide rapidly whether to form a coalition for taking advantage of the high reward. The first coalition that accepts the contract gets it and if multiple coalitions agree, one coalition is chosen at random. Agents that are only capable to perform a subset of the subtasks must propose or join a coalition. At each round, they can propose a coalition or accept to be part of one. Unlike in [82] where all the tasks are known in advance, in these works, a coalition is formed incrementally for each task. The order of the tasks may play an important role in the overall payoff of the agents.

If tasks arrive in a pattern, it may be efficient to form similar coalitions for similar tasks. Abdallah and Lesser assume the existence of a hierarchy of agents in [1]. When an agent gets a task for which it does not have the necessary resources, it can ask the agent above it in the hierarchy to take care of the task. If agents placed below it can solve subtasks of the task, the agent can decompose the task and assign it to the agents below in the hierarchy. Learning can be used to choose which agent can perform the task. Abdallah and Lesser show that learning allows for faster and better task assignments.

Another issue that can arise in the task allocation problem is the need to have overlapping coalitions, i.e., to have the possibility that agents are members of multiple coalitions. For example, an agent i may have a unique ability that is required to complete two tasks. If i is restricted to be a member of a single coalition, one task cannot be completed, which would be inefficient in the case where i has enough resources to help performing both tasks.

An example is in a transportation domain, if each task is to move an item between two points and a coalition is a set of vehicles that carry the item. Overlapping coalitions would model this problem as multiple items could be moved by the same truck [82].

The task allocation problem is in general a computationally hard problem: when agents are limited to perform a single task, the coalition problem resembles the set partitioning problem. When agents are able to perform multiple tasks, the allocation problem gets closer to the set covering problem. In both cases, these problems are $\mathcal{NP}$-complete [82]. A taxonomy is proposed to distinguish different complexity classes of the task allocation problems [50] based on three factors: (1) Is the same task likely to be offered again? (2) Does the multiagent system have more than enough/just enough/not enough resources to performing a set of tasks? (3) Is the reward intrinsic to the task, or does it only depend on the members performing the task? They show that some combinations of factors lead to polynomial problems, and other combinations have exponential complexity. Shehory and Kraus [82] restrict the set of possible coalitions by adding a constraint on the size of the coalition. This assumption is motivated by the fact that negotiation with a large number of partners becomes costly, and over a given size, a coalition of agents will not be able to get any benefit. In this case, the size of the set of possible coalitions is a constant, hence, the problem can be solved in polynomial time (possibly a high order polynomial though).

5.2 Electronic Marketplace

Coalition formation has also been used to model firms or agents in the electronic marketplace [7, 25, 52, 53, 77, 78, 89, 90]. The field originated from a paper by [89] where consumer agents can form a coalition (i.e., a buying group) to benefit from the quantity discount provided by sellers. From the point of view of a system designer, the problem is to form a CS, and each coalition is forming a buying group. Desirable property of the CS formed include to be Pareto optimal, i.e., no other CS should give more to a consumer without giving less to another one [7] and social welfare maximisation, which provides the greatest revenue to the buyers.

First, this problem can be modeled by coalitional games with non-transferable utility as in practice, a buyer may not pay another buyer to join a buying group (it is recognised that side payments could allow for more efficient outcomes though). It is the model studied by Asselin and Chaib-Draa [7]; their goal is to define protocols that finds a Pareto Optimal solutions, and they only propose a centralised solution.

One variant of the problem is to consider the cost of searching for other coalition members. For example, there is a cost to advertise the possibility to form a buying group, to look for partners, to negotiate price and payment [77]. The goal of the agents is to increase the size of the coalition so that the benefit from forming a coalition is worth the effort. The dilemma is about executing the task with the current configuration or starting a costly search to find new partners. The work of Sarne and Kraus analyses the equilibrium strategies.

Other variants include the problem introduced by Yamamoto and Sycara in [92]. Unlike in the original problem in [89], a buying group does not correspond to a particular item: each buyer agent can have a list of single items or a disjunction of items. In addition, sellers can bid discount prices to sell large volume of items. This allows a formulation that is closer to a combinatorial auction. The proposed solution assumes that each buying group is managed by an agent that has to solve the following two problems: (1) given the requests from the buyer agents, the manager agent chooses the sellers and buys the appropriate items, (2) the manager agent chooses the price paid by each buyer agent. To address the first problem, the proposed algorithm performs a greedy search. To answer the second problem, Yamamoto and Sycara use a surplus sharing rule that ensures a payoff distribution which is in the core. In [53], the agents can bid in combinatorial auctions: agents bid a reservation value for a bundle of items. This makes the problem even more complex since a winner determination problem has to be solved and a stable payoff distribution must be found. The mechanism design aspect of this problem can be found in [52]. Li and Sycara present an algorithm that computes an optimal coalition and a payoff division in the core in [53], but it is not guaranteed to be of polynomial complexity. Hence, they also present an approximation algorithm that is polynomial.

A mathematical model using first-order differential equations is presented in [51]. The model describes the dynamics of the coalitions and allows for computation of a steady state equilibrium. The paper shows that a steady state equilibrium always exists, and that it yields higher utility gain compared to the case where agents are buying on their own, or when leaving a coalition is not allowed. However, the outcome is not guaranteed to be Pareto Optimal.

Vassileva et al. address long term coalitions [17, 90]: in many other papers, a coalition is formed to complete a given task, and the coalition is disbanded when the task is accomplished. In contrast, the goal is to form a coalition of agents that will collaborate for a long period of time. The decision to leave a coalition and join a new one should also be a function of the trust put in another agents, i.e., the belief that they will have successful interaction in the future.

Another application of coalition formation in the context of an electronic marketplace is service oriented computing. A large number of services are offered on the Internet, at different prices and with different quality. Blankenburg et al. propose in [13] the use of service Request Agent that can request one (potentially) complex task and a Service Provider Agent that can provide a service. The latter can also, given a task and a set of service advertisements, compose services to form a plan that implements the task. The service requester agents only pay the service provider agent if the task is performed on time. The service provider agents must evaluate the risk involved in accepting a request. In addition, a service provider may be involved in more than one coalition, i.e., it can have multiple clients at the same time. Blankenburg et al. propose the use of fuzzy coalitions to allow agents to be member of multiple coalitions. The agents use a measure of risk from the finance literature, and they accept a proposal if the risk is below a threshold. To distribute the payoff, Blankenburg et al. define the kernel for their fuzzy coalition and use Stearns method to converge to a payoff distribution in the kernel.

5.3 Other Domains

Coalitions of agents have also been used in many other application domains, and we list some of them in the following. We start with an application for gathering information [44, 45]. An agent is associated with a local database, and to answer a query, an agent may require other agents. When the agents form a coalition, all agents in the coalition must cooperate: the members share some of their private data, e.g., dependency information. If an agent does not cooperate, it will not have access to some information schema that are available to members of the coalition. The coalition formation process assumes an utilitarian mechanism, and each agent tries to maximise its expected utility. The bilateral Shapley value is used to determine the payoff distribution in [44]. A kernel oriented solution is proposed in [45] for the same domain.

Coalitions have been used to track a moving target using a sensor network, a problem introduced in [39]. The problem is to ensure that at least three agents are sensing the target at the same time to perform triangulation. The problem becomes complex as the target is moving and sensors and communication can be faulty. In [84], the goal is for the agents to self-organise and form an appropriate coalition to track the target. The paper used a variant of the contract net protocol to negotiate a coalition that will be used throughout the tracking. Two valuation functions are studied (local and social utility) and different protocols are empirically tested. Soh et al. also solve a real-time tracking problem in [86]. An initiator agent starts the coalition formation process by contacting the neighboring agents that are most suitable for the particular task and engages in negotiation with each of them. Case based reasoning is used to choose the most promising negotiation protocol. In addition, reinforcement learning is used to estimate the utility of a coalition. The coalition formation process may or may not succeed.

In machine learning, it is known it is possible to combine the results of different classifier to increase the accuracy of the classification. [4] and [64] applied this idea in a coalition formation setting. For example, in the work of Plaza et al., agents can form committees (i.e., coalitions) to classify a new species of sponge. Each agent has its own expertise, a set of cases, and uses case-based reasoning for the classification problem. In their work, Plaza and Ontañón show how to decide when a committee is needed and how to select the agents to form a committee for a new species of sponge.

Coalitions of agents have been used in the context of distribution and planning of infrastructure for power transportation [23, 24, 65]. [65] model the trading process between firms that generate, transmit or distribute power using agents. Agents rank other agents by possible gains and send the Bilateral Shapley Value of the potential partner when it makes an offer. If both agents send requests to each other, it is beneficial for them to work together and they form a single entity. The process iterates until no further improvement is possible. In the power transmission domain, the problem is to decide whether or not to create a new line or a new plant, and if so, how to share the cost between the different parties involved. [23, 24] uses similar solutions as Yeung and Poon.

Coalitions of agents have been used in the context of planning and scheduling. For example, Pěchouček et al. tackle the problem of planning humanitarian relief operations in [62], and the problem of production planning in [63]. For the humanitarian relief operation scenario, different organisations can form coalitions to be more efficient and provide optimal help to the people. However, the different groups that have different capabilities can also have different goals; hence, they might not want to disclose all available information. In that context, the authors propose a formation of alliances: provided some public information, the agents seek to form groups of agents with the same kind of goals. These alliances can be viewed as long-term agreements between agents, and alliances define a partition of the agents. Unlike alliances, coalitions are viewed as short-term agreements to perform a specific task, and to reduce the search space, coalitions can form within an alliance. In case of impossibility of forming coalitions within an alliance, agents from different alliances can be used. The authors are interested in the amount of information agents have to disclose: when it sends a request, an agent may reveal private or semi-private information. This can occur when an agent asks an agent of a different alliance to perform a task (revealing that neither it nor its alliance can complete the task). An agent can also decide to disclose private information when it wants to inform other agents, for instance, when they form alliances.

In the context of production planning, instead of using a centralised planning approach, Pěchouček et al. want to use local coalition formation to execute tasks in an efficient manner. One requirement is that agents know their possible collaborators well in order to minimise the communication effort, e.g., agents have knowledge of the status of surrounding agents, so an agent may ask help from another agent if it knows the agent is not busy. Caillou et al. use the scenario of scheduling classes in a university [19], where a coalition is a schedule. This work considers non-transferable utility. Caillou et al. propose a protocol where a set of acceptable coalitions is passed from agents to agents, and each time, agents can add coalitions or remove coalitions that are not acceptable. The result of the protocol is a Pareto Optimal schedule. The protocol also considers re-using existing solutions to compute a solution to a modified problem (e.g., when a class is removed from the schedule, or a professor is coming, previous solutions of the problem can be used to accommodate these changes).

6 Issues

We now highlight issues that have emerged from the applications presented in Section 5. The protocols and algorithms we cited there present some solutions to these issues. Some additional issues remain unsolved, for example, dealing with agents that can enter and leave the environment at any time in an open, dynamic environment. None of the current protocols can handle these issues without re-starting computation, and only few approaches consider how to re-use the already computed solution [11, 19].

6.1 Stability and Dynamic Environments

Real-world scenarios often present dynamic environments. Agents can enter and leave the environment at any time, the characteristics of the agents may change with time, the knowledge of the agents about the other agents may change, etc.

The game-theoretic stability criteria are defined for a fixed population of agents and the introduction of a new agent in the environment requires significant computation to update a stable payoff distribution. For example, for the kernel, all the agents need to check whether any coalition that includes the new agent changes the value of the maximum surplus, which requires re-evaluating $O(2^n)$ coalitions. Given the complexity of the stability concept, one challenge that is faced by the multiagent community is to develop stability concepts that can be easily updated when an agent enters or leaves the environment.

In addition, if an agent drops during the negotiation, this may cause problems for the remaining agents. For example, a protocol that guarantees a kernel stable payoff distribution is shown not to be 'safe' when the population of agents is changing: if an agent i leaves the formation process without notifying other agents, the other agents may complete the protocol and find a solution to a situation that does not match the reality. Each time a new agent enters or leaves the population, a new process needs to be restarted [14].

In an open environments, manipulations will be impossible to detect: agents may use multiple identifiers (or false names) to pretend to be multiple agents, or the other way around, multiple agents may collude and pretend to be a single agents, or agents can hide some of their skills. Hence, it is important to propose solution concepts that are robust against such manipulations. We will come back later to some of the solution that have been proposed: the anonymity-proof core [93] and anonymity-proof Shapley value [60].

6.2 Uncertainty about Knowledge and Task

In real-world scenario, agents will be required to handle some uncertainty: agents may not know some tasks [14] or the value of some coalitions. In such cases, the agents play a different coalitional game that may reduce the payoff of some agents compared to the solution of the true game. Another example can be found in [46] where each agent knows the cost it incurs to perform a given task, but this information is considered private: an agent does not know the cost incurred by other agents and may only estimate these costs. This work assumes the presence of a trusted agent that plays the role of an auctioneer, and they assume that agents truthfully reveal their cost functions. In [43], Ketchpel proposes an auction-based protocol to distribute the payoff when there is uncertainty about the valuation of a coalition. In [86], agents also have uncertain and incomplete knowledge, and their approach is to use satisficing rather than optimal solution. Knowing the exact values of all coalitions is quite a strong assumption, and to relax it, [15] uses fuzzy sets to represent the coalition value. Another approach is to consider expected values of coalitions [20], and agents can have different expectations for a coalition value. This can also be the case when the value function is computationally hard to compute. In [75], computing a value for a coalition requires solving a version of the traveling salesman problem and approximations are used to solve that problem. In addition, when the agents do not use the same algorithm to compute the value of a coalition, not all agents may have the same value for each coalition.

6.3 Safety and Robustness

It is also important that the coalition formation process is robust. For instance, communication links may fail during the negotiation phase. Hence, some agents may miss some components of the negotiation stages. This possibility is studied in [14] for the KCA protocol [45]: coalition negotiations are not safe when some agents become unavailable (intentionally or otherwise). In particular, the payoff distribution is not guaranteed to be kernel-stable. [11] empirically studies the robustness of the use of a central algorithm introduced in [10]: the cost to compute a task allocation and payoff distribution in the core is polynomial, but it can still be expensive. In the case of agent failure, the computation needs to be repeated. Belmonte et al. propose an alternative payoff division model that avoids such a re-computation, but the solution is no longer guaranteed to be in the core, it is only close to the core. There is a trade-off between computational efficiency and the utility obtained by the agent. They conclude that when the number of agents is small, the loss of utility compared to the optimal is small; hence, the improvement of the computational efficiency can be justified. For a larger number of agents, however, the loss of utility cannot not justify the improvement in computational cost.

6.3.1 Protocol Manipulation

When agents send requests to search for members of a coalition or when they accept to form a coalition, the protocol may require disclosure of some private information [62]. When the agents reveal some of their information, the mechanism must ensure that there is no information asymmetry that can be exploited by some agents [12]. To protect a private value, some protocol [14] may allow the addition of a constant offset to the private value, as long as this addition does not impact the outcome of the negotiation.

Belmonte et al. study the effect of deception and manipulation of their model in [11]. They show that some agents can benefit from falsely reporting their cost. In some other approaches [14, 22], even if it is theoretically possible to manipulate the protocol, it is not possible in practice as the computational complexity required to ensure higher outcome to the malevolent agent is too high. For example, [22] show that manipulating marginal-contribution based value division scheme is $\mathcal{NP}$-hard (except when the valuation function has other properties, such as being convex).

Other possible protocol manipulations include hiding skills, using false names, colluding, etc. The traditional solution concepts can be vulnerable to false names and to collusion [93]. To address these problems, it is beneficial to define the valuation function in terms of the required skills instead of defining it over the agents: only skills, not agents, should be rewarded by the characteristic function. In that case, the solution concept is robust to false names, collusion, and their combination. But the agents can have incentive to hide skills. A straight, naive decomposition of the skills will increase the size of the characteristic function, and [94] propose a compact representation in this case.

6.4 Communication

While one purpose of better negotiation techniques may be to improve the quality of the outcome for the agents, other goals may include decreasing the time and the number of messages required to reach an agreement. For example, learning is used to decrease negotiation time in [85]. The motivation Lerman's work in [51] is to develop a coalition formation mechanism that has low communication and computation cost. In another work, the communication costs are included in the characteristic function [88].

The communication complexity of some protocols has been derived. For instance, the exponential protocol in [83] and the coalition algorithm for forming Bilateral Shapley Value Stable coalition in [44] have communication complexity of $O(n^2)$, the negotiation based protocol in [83] is $O(n^2 2^n)$, and it is $O(n^k)$ for the protocol in [82] (where k is the maximum size of a coalition). The goal of [66] is to analyse the communication complexity of computing the payoff of a player with different stability concepts: they find that it is $\Theta(n)$ when either the Shapley value, the nucleolus, or the core are used.

6.5 Scalability

When the population of heterogeneous agents is large, discovering the relevant agents to perform a task may be difficult. In addition, if all agents are involved in the coalition formation process, the cost in time and computation will be large. To alleviate this scalability issue, a hierarchy of agents can be used [1]. When an agent discovers a task that can be addressed by agents below this agent in the hierarchy, the agent picks the best of them to perform the task. If the agents below cannot perform the task, the agent passes the task to the agent above it in the hierarchy and the process repeats. The notion of clans [36] and congregations [18], where agents gather together for a long period have been proposed to restrict the search space by considering only a subset of the agents (see Section 6.6).

Another issue is the computational cost of the protocols for coalition formation. The nature of the problem is combinatorial: the size of the input representing a characteristic function in games with no externalities is at most $O(2^n)$ but with externalities, the input is even larger as it is $\Omega(n^{\frac{n}{2}})$. For large number of agents, it is not feasible to compute a payoff distribution that satisfies a stability criteria like the Shapley value or a kernel-stable payoff distribution. By restricting the size of the coalitions, kernel oriented coalition formation can be computed in polynomial time [45]. The use of bilateral Shapley value is also polynomial.

6.6 Long Term vs. Short Term

In general, a coalition is a short-lived entity that is "formed with a purpose in mind and dissolve when that need no longer exists, the coalition ceases to suit its designed purpose, or critical mass is lost as agents depart" [38]. It can be beneficial to consider the formation of long term coalitions, or the process of repeated coalition formation involving the same agents. The work by Vassileva and Breban explicitly study long term coalitions [90], and in particular the importance of trust in this content. Brooks and Durphee refer to a long term coalition as a congregation [18]. The purpose of a congregation is to reduce the number of candidates for a successful interaction: instead of searching the entire population, agents will only search in the congregation. The goal of a congregation is to gather agents, with similar or complementary expertise to perform well in an environment in the long run, which is not very different from a coalition. The only difference is that group rationality is not expected in a congregation. The notion of congregation is similar to the notion of clans [36]: agents gather not for a specific purpose, but for a long-term commitment. The notion of trust is paramount in the clans, and sharing information is seen as another way to improve performance.

6.7 Fairness

Stability does not necessarily imply fairness. For example, let us consider two CSs $\mathcal{S}$ and $\mathcal{T}$ with associated kernel-stable payoff distribution $x_{\mathcal{S}}$ and $x_{\mathcal{T}}$. Agents may have different preferences between the CSs. It may even be the case that there is no CS that is preferred by all agents. [2] consider games in partition function form (Section 2.8.1). They consider a process where, in turns, agents change coalition to improve their immediate payoff. They propose that the agents share the maximal social welfare, and the size of the share is proportional to the expected utility of the process. The payoff obtained is guaranteed to be at least as high as the expected utility. They claim that using the expected utility as a base of the payoff distribution provides some fairness as the expected utility can be seen as a global metric of an agent performance over the entire set of possible CSs. When agents are using the kernel as a stability concept, the payoff distribution depends on the CS which is formed by the agent. If some agents are in a singleton coaliton, this may not be fair compared to some agents that are in "good" coaliton. In [3], Airiau and Sen investigate the possibility of allowing transfer of utilities between members of different coalition in order to improve fairness.

6.8 Overlapping Coalitions

It is typically assumed that an agent belongs to a single coalition; however, there are some applications where agents can be members of multiple coalitions. As explained in the task allocation domain (see Section 5.1), the expertise of an agent may be required by different coalitions at the same time, and the agent can have enough resources to be part of two or more coalitions. In a traditional setting, the use of the same agent i by two coalitions $\mathcal{C}_1$ and $\mathcal{C}_2$ would require a merge of the two coalitions. This larger coalition U is potentially harder to manage, and a priori, there would not be much interaction between the agents in $\mathcal{C}_1$ and $\mathcal{C}_2$, except for agent i. Another application that requires the use of overlapping coalition is tracking targets using a sensor networks [91]. In this work, a coalition is defined for a target, and as agents can track multiple targets at the same time, they can be members of different coalitions.

The traditional stability concepts do not consider this issue. One possibility is for the agent to be considered as two different agents, but this representation is not satisfactory as it does not capture the real power of this agent. Shehory and Kraus propose a setting with overlapping coalition [82]: Each agent has a capacity, and performing a task may use only a fraction of the agent's capacity. Each time an agent commits to a task, the possible coalitions that can perform a given task can change. A mapping to a set covering problem allows to find the coalition. However, the study of the stability is not considered. Another approach is the use of fuzzy coalition [13]: agents can be members of a coalition with a certain degree that represents the risk associated with being in that coalition. Other work considers that the agents have different degree of membership, and their payoff depends on this degree [8, 55, 59]. The protocols in [50] also allow overlapping coalitions.

More recently, [21] have studied the notion of the core in overlapping coalition formation. In their model, each agent has one resource and the agent contributes a fraction of that resource to each coalition it participates in. The

valuation function v is then $[0,1]^n \to \mathbb{R}$. A CS is no longer a partition of the agents: a CS S is a finite list of vectors, one for each 'partial' coalition, i.e., $S = (r^1, \ldots, r^k)$. The size of S is the number of coalitions, i.e., k. The support of $r^{\mathcal{C}} \in S$ (i.e., the set of indices $i \in N$ such that $r_i^{\mathcal{C}} \neq 0$) is the set of agents forming coalition $\mathcal{C}$. For all $i \in N$ and all coalition $\mathcal{C} \in S$, $r_i^{\mathcal{C}} \in [0,1]^n$ represents the fraction of resource that agent i contributes to coalition $\mathcal{C}$; hence, $\sum_{\mathcal{C} \in S} r_i^{\mathcal{C}} \leq 1$ (i.e., agent i cannot contributes more than 100% of its resource). A payoff distribution for a CS S of size k is defined by a k-tuple $x = (x^1, \ldots, x^k)$ where $x^{\mathcal{C}}$ is the payoff distribution that the agents obtain for coalition $\mathcal{C}$. If an agent is not in the coalition, it must not receive any payoff for this coalition, hence $(r_i^{\mathcal{C}} = 0) \Rightarrow (x_i^{\mathcal{C}} = 0)$. The total payoff of agent i is the sum of its payoff over all coalitions $p_i(CS, x) = \sum_{\mathcal{C}=1}^{k} x_i^{\mathcal{C}}$. The efficiency criterion becomes $\forall r^{\mathcal{C}} \in S$, $\sum_{i \in N} x_i^{\mathcal{C}} = v(r^{\mathcal{C}})$. An imputation is an efficient payoff distribution that is also individually rational. We denote by $I(S)$ the set of all imputations for the CS S. We are now ready to define the overlapping core: a pair (S, x) is in the *overlapping core* when for any set of agents $\mathcal{C} \subseteq N$, any CS S', any imputation $y \in I(S')$ we have $\exists i \in N$, $p_i(S', y) \leq p_i(S, x)$. The work in [21] provides a characterisation of the core under some (mild) conditions for the utility function.

7 Conclusion

In this tutorial, we reviewed cooperative game theory. We introduced the two main models: the TU and the NTU games. For the case of TU games, the concept of the core is central. It has a natural definition to promote stability. Unfortunately, the core of a game is sometimes empty, and different stability concepts have been introduced to provide some solutions. In particular, we introduced the kernel and the nucleolus. We also introduced the Shapley value, which promotes a fair (and not necessarily stable) payoff distribution. A priori, none of the stability concepts introduced has a clear advantage over another. A given application may favour one solution concepts. Note that more solution concepts have been introduced (for example the bargaining set or some refinement of the core such as the ϵ-core or the least core). We also briefly introduced NTU games, and we only introduced the notion of the core. Other stability concepts can also be introduced for these games.

The remaining part of the tutorial showed that the models offered by cooperative game theory may not be sufficient to answer some potential real world problems, and we surveryed some issues raised in the multiagent systems literature. Not all the issues have been successfully resolved and many problems remain open.

References

[1] Sherief Abdallah and Victor Lesser. Organization-based cooperative coalition formation. In *IAT 04*, 2004.

[2] Stéphane Airiau and Sandip Sen. A fair payoff distribution for myopic rational agents. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-09)*, pages 1305–1306, May 2009.

[3] Stéphane Airiau and Sandip Sen. On the stability of an optimal coalition structure. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, August 2010. To appear.

[4] Samir Aknine and Philippe Caillou. Agreements without disagreements. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004*, pages 3–7, August 2004.

[5] Samir Aknine, Suzanne Pinson, and Melvin F. Shakun. A multi-agent coalition formation method based on preference models. *Group Decision and Negotiation*, 13(6):513–538, November 2004.

[6] Samir Aknine and Onn Shehory. Coalition Formation: Concessions, Task Relationships and Complexity Reduction. *ArXiv Computer Science e-prints*, February 2005.

[7] Frederick Asselin and Brahim Chaib-draa. Performance of software agents in non-transferable payoff group buying. *Journal of Experimental and Theoretical Artificial Intelligence*, 18(1):1–32, mar 2006.

[8] J.-P. Aubin. *Mathematical Methods of Game and Economic Theory*. North-Holland, 1979.

[9] Robert J. Aumann and Jacques H Drèze. Cooperative games with coalition structures. *International Journal of Game Theory*, 3(4):217–237, 1974.

[10] María-Victoria Belmonte, Ricardo Conejo, José-Luis Pérez de-la Cruz, and Francisco Triguero Ruiz. A stable and feasible payoff division for coalition formation in a class of task oriented domains. In *ATAL '01: Revised Papers from the 8th International Workshop on Intelligent Agents VIII*, pages 324–334, London, UK, 2002. Springer-Verlag.

[11] María-Victoria Belmonte, Ricardo Conejo, José-Luis Pérez de-la Cruz, and Francisco Triguero Ruiz. A robust deception-free coalition formation model. In *Proceedings of the 2004 ACM symposium on Applied computing (SAC '04)*, pages 469–473, New York, NY, USA, 2004. ACM Press.

[12] Bastian Blankenburg, Rajdeep K. Dash, Sarvapali D. Ramchurn, Matthias Klusch, and Nicholas R. Jennings. Trusted kernel-based coalition formation. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 989–996, New York, NY, USA, 2005. ACM Press.

[13] Bastian Blankenburg, Minghua He, Matthias Klusch, and Nicholas R. Jennings. Risk-bounded formation of fuzzy coalitions among service agents. In *Proceedings of 10th International Workshop on Cooperative Information Agents*, 2006.

[14] Bastian Blankenburg and Matthias Klusch. On safe kernel stable coalition forming among agents. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems*, 2004.

[15] Bastian Blankenburg, Matthias Klusch, and Onn Shehory. Fuzzy kernel-stable coalitions between rational agents. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems (AAMAS-03)*. ACM Press, 2003.

[16] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

[17] Silvia Breban and Julita Vassileva. A coalition formation mechanism based on inter-agent trust relationships. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 306–307, New York, NY, USA, 2002. ACM Press.

[18] Christopher H. Brooks and Edmund H. Durfee. Congregation formation in multiagent systems. *Autonomous Agents and Multi-Agent Systems*, 7(1-2):145–170, 2003.

[19] Philippe Caillou, Samir Aknine, and Suzanne Pinson. A multi-agent method for forming and dynamic restructuring of pareto optimal coalitions. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 1074–1081, New York, NY, USA, 2002. ACM Press.

[20] Georgios Chalkiadakis and Craig Boutilier. Bayesian reinforcement learning for coalition formation under uncertainty. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems(AAMAS'04)*, 2004.

[21] Georgios Chalkiadakis, Edith Elkind, Evangelos Markakis, and Nicholas R. Jennings. Overlapping coalition formation. In *Proceeedings of the 4th International Workshop on Internet and Network Economics (WINE2008)*, pages 307–321, 2008.

[22] Vincent Conitzer and Tuomas Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-04)*, pages 219–225, 2004.

[23] Javier Contreras, Matthias Klusch, Onn Shehory, and Felix Wu. Coalition formation in a power transmission planning environment. In *Proceeding of the Second International Conference on Practical Applications of Multi-Agent Systems, PAAM97*, April 1997.

[24] Javier Contreras, Matthias Klusch, and Jerome Yen. Multi-agent coalition formation in power transmission planning: a bilateral shapley value approach. In *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems*, 1998.

[25] David Cornforth, Michael Kirley, and Terry Bossomaier. Agent heterogeneity and coalition formation: Investigating market-based cooperative problem solving. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems(AAMAS'04)*, 2004.

[26] Viet Dung Dang and Nicholas R. Jennings. Generating coalition structures with finite bound from the optimal guarantees. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems(AAMAS'04)*, 2004.

[27] M. Davis and M. Maschler. The kernel of a cooperative game. *Naval Research Logistics Quarterly*, 12, 1965.

[28] Bart de Keijzer, Tomas Klos, and Yingqian Zhang. Enumeration and exact design of weighted voting games. In *Proc. of the 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2010)*, pages 391–398, 2010.

[29] Xiaotie Deng, Qizhi Fang, and Xiaoxun Sun. Finding nucleolus of flow game. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 124–131, New York, NY, USA, 2006. ACM Press.

[30] Xiaotie Deng and C H Papadimitriou. On the complexity of cooperative solution concetps. *Mathematical Operation Research*, 19(2):257–266, 1994.

[31] Tone Dieckmann and Ulrich Schwalbe. Dynamic coalition formation and the core. *Journal of Economic Behavior & Organization*, 49(3):363–380, November 2002.

[32] Edith Elkind, Georgios Chalkiadakis, and Nicholas R. Jennings. Coalition structures in weighted voting games. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08)*, pages 393–397, 2008.

[33] Edith Elkind, Leslie Ann Goldberg, Paul Goldberg, and Michael Wooldridge. Computational complexity of weighted threshold games. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI-07)*, pages 718–723, 2007.

[34] Edith Elkind and Michael Wooldridge. Hedonic coalition nets. In *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-09)*, pages 417–424, May 2009.

[35] Donald B. Gillies. *Some theorems on n-person games*. PhD thesis, Department of Mathematics, Princeton University, Princeton, N.J., 1953.

[36] Nathan Griffiths and Michael Luck. Coalition formation through motivation and trust. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems (AAMAS03)*, New York, NY, USA, 2003. ACM Press.

[37] Sergiu Hart and Mordecai Kurz. Endogenous formation of coalitions. *Econometrica*, 51(4), July 1983.

[38] Bryan Horling and Victor Lesser. A survey of multi-agent organizational paradigms. *The Knowledge Engineering Review*, 19:281–316, 2004.

[39] Bryan Horling, Régis Vincent, Roger Mailler, Jiaying Shen, Raphen Becker, Kyle Rawlins, and Victor Lesser. Distributed sensor network for real time tracking. In *Proceedings of the fifth international conference on Autonomous agents(Agents'01)*, pages 417–424, New York, NY, USA, 2001. ACM Press.

[40] Samuel Ieong and Yoav Shoham. Marginal contribution nets: a compact representation scheme for coalitional games. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 193–202, New York, NY, USA, 2005. ACM Press.

[41] James P. Kahan and Amnon Rapoport. *Theories of Coalition Formation*. Lawrence Erlbaum Associates, Publishers, 1984.

[42] Steven P. Ketchpel. The formation of coalitions among self-interested agents. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 414–419, August 1994.

[43] Steven P. Ketchpel. Forming coalitions in the face of uncertain rewards. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 414–419, August 1994.

[44] Matthias Klusch and Onn Shehory. Coalition formation among rational information agents. In Rudy van Hoe, editor, *Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, Eindhoven, The Netherlands, 1996.

[45] Matthias Klusch and Onn Shehory. A polynomial kernel-oriented coalition algorithm for rational information agents. In *Proceedings of the Second International Conference on Multi-Agent Systems*, pages 157 – 164. AAAI Press, December 1996.

[46] Sarit Kraus, Onn Shehory, and Gilad Taase. Coalition formation with uncertain heterogeneous information. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 1–8. ACM Press, 2003.

[47] Sarit Kraus, Onn Shehory, and Gilad Taase. The advantages of compromising in coalition formation with incomplete information. In *Proceedings of the third International Joint Conference on Autonomous Agents and Multiagent Systems(AAMAS'04)*, 2004.

[48] Jeroen Kuipers, Ulrich Faigle, and Walter Kern. On the computation of the nucleolus of a cooperative game. *International Journal of Game Theory*, 30(1):79–98, 2001.

[49] Kate S. Larson and Tuomas W. Sandholm. Anytime coalition structure generation: an average case study. *Journal of Experimental & Theoretical Artificial Intelligence*, 12(1):23–42, 2000.

[50] Hoong Chuin Lau and Lei Zhang. Task allocation via multi-agent coalition formation: Taxonomy, algorithms and complexity. In *15th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2003)*, pages 346–350, November 2003.

[51] Katia Lerman and Onn Shehory. Coalition formation for large-scale electronic markets. In *Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*. IEEE Computer Society, July 2000.

[52] Cuihong Li, Uday Rajan, Shuchi Chawla, and Katia Sycara. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 68–77, New York, NY, USA, 2003. ACM Press.

[53] Cuihong Li and Katia Sycara. Algorithm for combinatorial coalition formation and payoff division in an electronic marketplace. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems(AAMAS '02)*, pages 120–127, New York, NY, USA, 2002. ACM Press.

[54] Efrat Manisterski, Esther David, Sarit Kraus, and Nicholas R. Jennings. Forming efficient agent groups for completing complex tasks. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS-06*, May 2006.

[55] M. Mares. Fuzzy cooperative games: cooperation with vague expectations. *Studies in fuzziness and soft computing*, 72, 2001.

[56] Tomasz Michalak, Andrew Dowell, Peter McBurney, and Michael Wooldridge. Optimal coalition structure generation in partition function games. In *Proceeding of the 2008 conference on ECAI 2008*, pages 388–392, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press.

[57] Tomasz Michalak, Talal Rahwan, Jacek Sroka, Andrew Dowell, Michael Wooldridge, Peter McBurney, and Nicholas R. Jennings. On representing coalitional games with externalities. In *Proceedings of the 10th ACM conference on Electronic Commerce 09 (EC'09)*, 2009.

[58] Roger B. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, 2:225–229, 1977.

[59] I. Nishizaki and M. Sakawa. Masatoshi. Fuzzy and multiobjective games for conflict resolution. *Studies in Fuzziness and Soft Computing*, 64, 2001.

[60] Naoki Ohta, Vincent Conitzer, Yasufumi Satoh, Atsushi Iwasaki, and Makoto Yokoo. Anonimity-proof shapley value: Extending shapley value for coalitional games in open environments. In *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS-09)*, May 2009.

[61] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

[62] Michal Pěchouček, Vladimír Mařík, and Jaroslav Bárta. A knowledge-based approach to coalition formation. *IEEE Intelligent Systems*, 17(3):17–25, 2002.

[63] Michal Pěchouček, Vladimír Mařík, and Olga čStépánková. Coalition formation in manufacturing multi-agent systems. In *Proceedings of the 11th International Workshop on Database and Expert Systems Applications (DEXA'00)*, 2000.

[64] Enric Plaza and Santiago Ontañón. Learning collaboration strategies for committees of learning agents. *Journal of Autonomous Agents and Multi-Agent Systems*, 13(3):429–461, 2006.

[65] Ada S. Y. Poon, Chris S. K. Yeung, and Felix F. Wu. Game theoretical multi-agent modeling of coalition formation for multilateral trades. *IEEE Transactions on Power Systems*, 4(3), August 1999.

[66] Ariel D. Procaccia and Jeffrey S. Rosenschein. The communication complexity of coalition formation among autonomous agents. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS-06*, May 2006.

[67] Talal Rahwan and Nicholas R. Jennings. An algorithm for distributing coalitional value calculations among cooperating agents. *Artificial Intelligence*, 171(8-9):535–567, 2007.

[68] Talal Rahwan and Nicholas R. Jennings. Coalition structure generation: Dynamic programming meets anytime optimization. In *Proceedings of the 23rd conference on artificial intelligence (AAAI-08)*, pages 156–161, 2008.

[69] Talal Rahwan and Nicholas R. Jennings. An improved dynamic programming algorithm for coalition structure generation. In *Proceedings of the 7th international conference on Autonomous Agents and Multi-Agent Systems (AAMAS-08)*, 2008.

[70] Talal Rahwan, Tomasz Michalak, Nicholas R. Jennings, Michael Wooldridge, and Peter McBurney. Coalition structure generation in multi-agent systems with positive and negative externalities. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-09)*, 2009.

[71] Talal Rahwan, Sarvapali D. Ramchurn, Viet Dung Dang, Andrea Giovannucci, and Nicholas R. Jennings. Anytime optimal coalition structure generation. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, pages 1184–1190, 2007.

[72] Talal Rahwan, Sarvapali D. Ramchurn, Viet Dung Dang, and Nicholas R. Jennings. Near-optimal anytime coalition structure generation. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 2365–2371, January 2007.

[73] Talal Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, and Andrea Giovannucci. An anytime algorithm for optimal coalition structure generation. *Journal of Artificial Intelligence Research*, 34:521–567, 2009.

[74] Debraj Ray and Rajiv Vohra. A theory of endogenous coalition structures. *Games and Economic Behavior*, 26:286–336, 1999.

[75] Tuomas Sandholm and Victor R. Lesser. Coalitions among computationally bounded agents. *AI Journal*, 94(1–2):99–137, 1997.

[76] Tuomas W. Sandholm, Kate S. Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999.

[77] David Sarne and Sarit Kraus. The search for coalition formation in costly environments. In *Cooperative Information Agents VII, 7th International Workshop, CIA 2003, Helsinki, Finland, August 27-29, 2003, Proceedings*, volume 2782 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, August 2003.

[78] David Sarne and Sarit Kraus. Cooperative exploration in the electronic marketplace. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, pages 158–163. AAAI Press AAAI Press / The MIT Press, 2005.

[79] D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal of applied mathematics*, 17, 1969.

[80] Sandip Sen and Partha Sarathi Dutta. Searching for optimal coalition structures. In *ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, page 287, Washington, DC, USA, 2000. IEEE Computer Society.

[81] L.S. Shapley. A value for n-person games. In H. Kuhn and A.W. Tucker, editors, *Contributions to the Theory of Games*, volume 2. Princeton University Press, Princeton, NJ, 1953.

[82] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, May 1998.

[83] Onn Shehory and Sarit Kraus. Feasible formation of coalitions among autonomous agents in nonsuperadditve environments. *Computational Intelligence*, 15:218–251, 1999.

[84] Mark Sims, Claudia V. Goldman, and Victor Lesser. Self-organization through bottom-up coalition formation. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems (AAMAS03)*, pages 867–874, New York, NY, USA, 2003. ACM Press.

[85] Leen-Kiat Soh and Costas Tsatsoulis. Satisficing coalition formation among agents. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 1062–1063, New York, NY, USA, 2002. ACM Press.

[86] Leen-Kiat Soh, Costas Tsatsoulis, and Huseyin Sevay. A satisficing, negotiated, and learning coalition formation architecture. *Distributed Sensor Networks: A Multiagent Perspective*, pages 109–138, 2003.

[87] Richard Edwin Stearns. Convergent transfer schemes for n-person games. *Transactions of the American Mathematical Society*, 134(3):449–459, December 1968.

[88] Fernando Tohmé and Tuomas Sandholm. Coalition formation processes with belief revision among bounded-rational self-interested agents. *Journal of Logic and Computation*, 9(6):793–815, 1999.

[89] Maksim Tsvetovat, Katia P. Sycara, Yian Chen, and James Ying. Customer coalitions in electronic markets. In *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Electronic Commerce Systems (includes revised papers from AMEC 2000 Workshop)*, pages 121–138, London, UK, 2001. Springer-Verlag.

[90] Julita Vassileva, Silvia Breban, and Michael Horsch. Agent reasoning mechanism for making long-term coalitions based on decision making and trust. *Computational Intelligence*, 18(4):583–595, 2002.

[91] Alex Rogers Viet Dung Dang, Rajdeep K. Dash and Nicholas R. Jennings. Overlapping coalition formation for efficient data fusion in multi-sensor networks. In *Proceedings of the Twenty-First Conference on Artificial Intelligence (AAAI-06)*, pages 635–640, 2007.

[92] Junichi Yamamoto and Katia Sycara. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 576–583, New York, NY, USA, 2001. ACM Press.

[93] Makoto Yokoo, Vincent Conitzer, Tuomas Sandholm, Naoki Ohta, and Atsushi Iwasaki. Coalitional games in open anonymous environments. In *Proceedings of the Twentieth National Conference on Artificial Intelligence*, pages 509–515. AAAI Press AAAI Press / The MIT Press, 2005.

[94] Makoto Yokoo, Vincent Conitzer, Tuomas Sandholm, Naoki Ohta, and Atsushi Iwasaki. A compact representation scheme for coalitional games in open anonymous environments. In *Proceedings of the Twenty First National Conference on Artificial Intelligence*, pages –. AAAI Press AAAI Press / The MIT Press, 2006.

[95] H. P. Young. Monotonic solutions of cooperative games. *International Journal of Game Theory*, 14:65–72, 1985.