Forming Coalitions

Cooperative Game Theory

Cooperative Games

- **n** agents (n > 2)
- **Ag = {1, …, n}** - set of all agents
- Refer to a subset of **Ag** as a **coalition**
- A **coalition** is a set of agents, which may or may not work together
- **Grand coalition** – coalition consisting of all agents **C = Ag**
- **Singleton coalition** – coalition containing one agent
- A coalition have the ability to obtain a certain utility, which can be shared among coalition members

Cooperative Games

- A **cooperative game** (**coalitional game**) is a pair:
- (**Ag, v**), where **Ag** is a set of agents and
- $\mathbf{v: 2^{Ag} \rightarrow R}$ is called the **characteristic function** of the game
- The characteristic function assigns to every possible coalition the pay-off that may be distributed between the members of that coalition

Example

- **v(C) = k** → coalition **C** can cooperate in such a way that they will obtain utility **k**, which may then be distributed among team members
- Coalition members have to agree among themselves how to divide the utility
- **v(C)** is the largest value that **C** could obtain by cooperating, and we are not concerned with how they cooperate

Simple Coalitional Game

- Game where the value of any coalition is either **0** (coalition is **losing**) or **1** (coalition is **winning**)

Questions

- Which coalitions might be formed by rational agents ?
- How the pay-off received by a coalition might be reasonably divided between the members of a coalition ?

Which Coalition Should I Join?

- Is a coalition **stable** ? that is
- Is it rational for all members of coalition to stay with the coalition, or could they benefit by defecting from it ?
- Stability is a *necessary* but not *sufficient* condition for coalitions to form
- Consider the stability of the grand coalition

The Core

- The **core** of a coalitional game is the set of *feasible* distributions of payoff to members of a coalition that *no* sub-coalition can reasonably object to
- An outcome **x** for a coalition **C** in a game (**Ag, v**) is a distribution of **C**'s utility to members of **C**, $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$, which is both **feasible** for **C** (**C** could obtain the pay-off indicated in x) and **efficient** (all available utility is allocated)

The Core

- An outcome **$\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$** for coalition **$\mathbf{C} = \{1, \ldots, k\}$** must satisfy the property:
- $\sum_{i \in C} x_i$ = v(C)

Example

- Game with Ag = {1, 2}, such that:
- v({1}) = 5, v({2}) = 5, and v({1,2}) = 20
- The possible outcomes are:
- (20,0), (19,1), (18,2), …, (0,20)
- Outcome (20,0) → agent 1 gets the entire collective value, while agent 2 gets nothing

Objections

- A coalition **C objects** to an outcome for the grand coalition, if there is some outcome *for C* that makes *all members of C* strictly better off
- **C $\subseteq$ Ag objects** to an outcome $\mathbf{(x_1, \ldots, x_n)}$ for the grand coalition if there is some outcome $\mathbf{(x_1', \ldots, x_k')}$ for **C** such that:
- $\mathbf{x_i' > x_i} \;\; \forall\, i \in C$

Observation

- An outcome is not going to happen if some coalition objects to it – they would do better by defecting and working on their own

The Core

- The **core** is the set of outcomes for the *grand coalition* to which *no* coalition objects
- If the core is *non-empty*, then the *grand coalition is stable*, since nobody can benefit from defection
- Asking:
- *Is the grand coalition stable ?*
- is the same as asking:
- *Is the core non-empty ?*

Problems with the Core

- Is the core non-empty in the previous example ?
- Enumerate all outcomes for the grand coalition, and ask whether any coalition objects
- (20,0) – agent 2 objects, since it can do better on its own: v({2}) = 5
- (19,1), (18,2) – agent 2 continues to object

Problems with the Core

- (15,5) – agent 2 ceases to have any objection, since it cannot do better on its own
- Agent 2 would probably argue that outcome (15,5) is unfair
- (14,6), …, (5,15) – no objection
- (4,16) – agent 1 starts to object all the way to (0,20)
- The core contains all the outcomes between (15,5) and (5,15)

Shapley Value

- The *Shapley value* is best known attempt to define how to divide benefits of cooperation fairly
- It does this by taking into account *how much an agent contributes*

Definitions

- Let **μ_i(C)** be the amount that **i** adds by joining **C** $\subseteq$ **Ag \ {i}** - the **marginal contribution of i to C**
- **$\mu_i(C) = v(C \cup \{i\}) – v(C)$**
- If **$\mu_i(C) = v(\{i\})$**, then there is no added value from **i** joining **C**, since the amount **i** adds is exactly what **i** would earn on its own
- **sh_i** is the value that agent **i** $\in$ **Ag** is given in the game **(Ag, v)**

Symmetry

- Agents that make the same contribution should get the same pay-off
- The amount an agent gets should only depend on their contribution, and not on their name
- **i** and **j** are **interchangeable** if $\mu_i(C) = \mu_j(C)$, for every $C \subseteq Ag \setminus \{i,j\}$
- If **i** and **j** are interchangeable then $sh_i = sh_j$

Dummy Player

- Agents that never have any synergy with any coalition should get only what they can earn on their own
- $\mathbf{i} \in \mathbf{Ag}$ is a **dummy player** if $\mu_i(\mathbf{C}) = \mathbf{v}(\{\mathbf{i}\})$, for every $\mathbf{C} \subseteq \mathbf{Ag} \setminus \{\mathbf{i}\}$
- A dummy is a player that only adds to a coalition the value that it could obtain on its own
- If **i** is a dummy player, then $\mathbf{sh_i} = \mathbf{v}(\{\mathbf{i}\})$

Additivity

- If you combine two games, then the value that a player gets should be the sum of the values it gets in the individual games
- A player does not gain or lose by playing more than once

Additivity

- Let (**Ag,** $\mathbf{v^1}$) and (**Ag,** $\mathbf{v^2}$) be two games containing the same set of agents
- Let **i** $\in$ **Ag** be one of the agents
- Let $\mathbf{sh_i^1}$ and $\mathbf{sh_i^2}$ be the value player **i** receives in the two games, respectively
- Let (**Ag,** $\mathbf{v^{1+2}}$) be the game such that:
- $\mathbf{v^{1+2}(C) = v^1(C) + v^2(C)}$

Additivity

- The value $\mathbf{sh_i^{1+2}}$ of player **i** in the game ($\mathbf{Ag, v^{1+2}}$) should be $\mathbf{sh_i^1 + sh_i^2}$

Modular Representations - Induced Subgraphs

- A characteristic function is defined by an undirected, weighted graph, in which nodes in the graph are the members of **Ag**
- The edges are labelled with integer values
- Let $\mathbf{w}_{\mathbf{i,j}}$ be the weight of the edge from **i** to **j** in the graph

Induced Subgraphs

- To compute the value of a coalition **C** $\subseteq$ **Ag**, sum all the edges in the graph whose components are all contained in **C** – the value of a coalition **C** $\subseteq$ **Ag** is the weight of the **subgraph induced by C**

$$\nu(C) = \sum_{\{i,j\} \subseteq Ag} w_{i,j}$$

Example

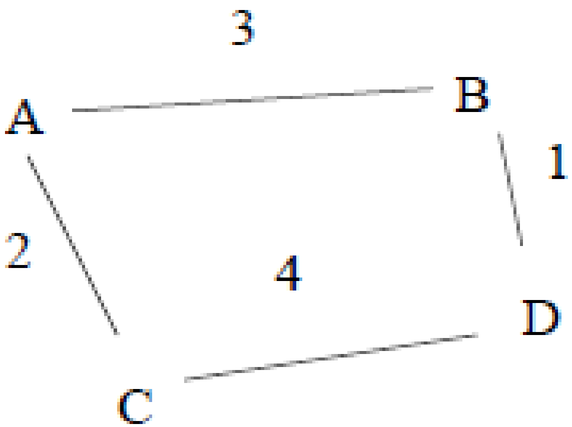

the original graph defining v

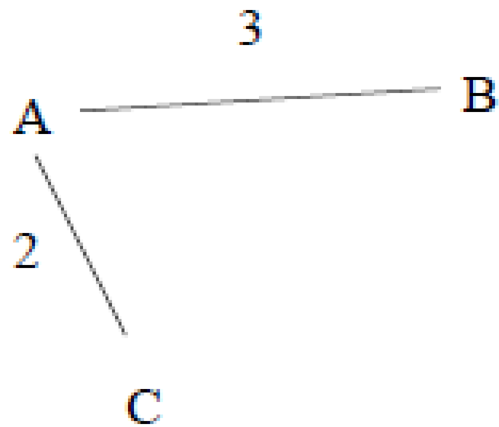

subgraph induced by {A,B,C}
giving v({A,B,C}) = 3 + 2 = 5

Observation

- Using an adjacency matrix to represent the graph requires **|Ag|^2** entries
- This is not a complete representation – there are characteristic functions **v** that cannot be represented using such a graph

Modular Representations – Marginal Contribution Nets

- Represent the characteristic function of a game as a set of rules of the form:
- **pattern → value**
- The pattern is a conjunction of agents
- A rule applies to a group of agents **C** if **C** is a superset of the agents in the pattern

Example

- A rule with pattern $1 \wedge 3$ would apply to the coalitions {1, 3}, {1, 3, 5}, but not to the coalition {1} or {8, 12}

Definitions

- If $\phi \rightarrow$ **x** is a rule and **C** is a coalition, then **C |=** ϕ means that the rule applies to coalition **C**
- Given a set of rules **rs**, denote by $\mathbf{rs_C}$ the subset of **rs** that apply to coalition **C**
- $\mathbf{rs_C = \{\phi \rightarrow x \in rs \mid C \models \phi\}}$

Definitions

- Given a set of rules **rs**, we compute the value of a given coalition **C** by summing the values of all the rules that apply to the coalition
- The characteristic function $\mathbf{v_{rs}}$ associated with the set of rules **rs** is defined by:

$$v_{rs}(C) = \sum_{\phi \rightarrow x \in rs_C} x$$

Example 1

- Rules set rs1 with the rules:
- $a \wedge b \rightarrow 5$
- $b \rightarrow 2$
- We have:
- $v_{rs1}(\{a\}) = 0$
- $v_{rs1}(\{b\}) = 2$
- $v_{rs1}(\{a,b\}) = 7$

Example 2

- Rules set rs2 with the rules:
- $a \wedge b \rightarrow 5$
- $b \rightarrow 2$
- $c \rightarrow 4$
- $b \wedge \sim c \rightarrow -2$

- $v_{rs2}(\{a\}) = 0$ (no rules apply)
- $v_{rs2}(\{b\}) = 0$ (second and fourth rules)
- $v_{rs2}(\{c\}) = 4$ (third rule)
- $v_{rs2}(\{a,b\}) = 5$ (first, second, and fourth rules)
- $v_{rs2}(\{a,c\}) = 4$ (third rule)
- $v_{rs2}(\{b,c\}) = 6$ (second and third rule)

Simple Coalitional Games

- Game in which every coalition gets a value of either **1** (winning) or **0** (losing)
- Simple Coalitional Games model **yes/no** voting systems

Model Simple Coalitional Games

- Pair **Y = {Ag, W}**
- **Ag = {1, …, n}** set of agents or voters
- $\mathbf{W} \subseteq \mathbf{2^{Ag}}$ set of winning coalitions – if $\mathbf{C} \in \mathbf{W}$, then **C** would be able to determine the outcome (**yes/no**) to the question at hand, should they collectively choose to

Scenarios

- **Non-triviality** – there are some winning coalitions, but not all coalitions are winning - $\emptyset \subset \mathbf{W} \subset \mathbf{2}^{\mathbf{Ag}}$
- **Monotonicity** – if **C** wins, then every superset of **C** also wins – if $\mathbf{C_1} \subseteq \mathbf{C_2}$ and $\mathbf{C_1} \in \mathbf{W}$ then $\mathbf{C_2} \in \mathbf{W}$
- **Zero-sum** – if a coalition **C** wins, then the agents outside **C** do not win – if $\mathbf{C} \in \mathbf{W}$ then $\mathbf{Ag} \setminus \mathbf{C} \notin \mathbf{W}$

Weighted Voting Games

- For each agent **i** ∈ **Ag**, we define a weight $\mathbf{w_i}$ and an overall quota **q**
- A coalition **C** is winning if the sum of their weights exceeds the quota:

$$\nu(C) = \begin{cases} 1 & \text{if } \sum_{i \in C} w_i \geq q \\ 0 & \text{otherwise.} \end{cases}$$

Notation

- A weighted voting game with players 1,…,n, having weights $w_1,\ldots,w_n$ and quota q is written as:
- **$(q;w_1,\ldots,w_n)$**

Observation

- Simple majority voting is a special case of weighted voting games
- Players are the people voting, and each player has weight $w_i = 1$
- The quota q is:
- $q = (|Ag| + 1)/2$

Example 1

- Consider a weighted voting game, with agents Ag = {1, 2, 3}
- (100; 99, 99, 1)
- Intuition suggests that since agents 1 and 2 have weights nearly 100 times larger than that of agent 3, then they must be more powerful
- Any coalition containing at least two players is winning – the three agents are interchangeable

Example 2

- Sometimes, the fact that a voter has a non-zero power is meaningless
- Consider a weighted voting game
- (10; 6, 4, 2)
- The final player is never able to influence a decision – the only winning coalitions are those containing at least the first two players, and adding the final player to a coalition never changes the coalition status

Example 3

- Add a new voter, keeping the quota unchanged
- (10; 6, 4, 2, 8)
- The coalition of the final two players is winning
- This does not prevent the first two from also being winning
- This game is not zero-sum

Observation

- The core of a weighted voting game is non-empty iff there is an agent present in every winning coalition
- To check whether **i** is present in every winning coalition, use a list **C** of all the agents, except **i**, that have positive weights
- If **i** is supposed to be present in all winning coalitions, then **C** must be losing, and adding **i** to **C** must make them winning

Coalitional Games with Goals

- Qualitative coalitional games (QCG) are a type of coalitional game in which each agent has a set of goals, and wants one of them to be achieved (but it doesn’t care which)
- Agents cooperate in QCG to achieve mutually satisfying sets of goals

Qualitative Coalitional Games

- There is an overall set of possible goals **G**
- Every agent $\mathbf{i} \in \mathbf{Ag}$ has some set of goals $\mathbf{G_i} \subseteq \mathbf{G}$ associated with it
- The agent wants one of these goals to be achieved, but does not care which one
- Every coalition **C** has a set of choices **V(C)** associated with it, representing the different ways that the coalition **C** could choose to cooperate

Qualitative Coalitional Games

- The characteristic function for QCG is:

$$V : 2^{Ag} \rightarrow 2^{2^{G}}$$

Qualitative Coalitional Games

- Suppose that a set of goals $\mathbf{G'} \subseteq \mathbf{G}$ is achieved
- **G'** is said to **satisfy** an agent **i** if $\mathbf{G_i} \cap \mathbf{G'} \neq \emptyset$, that is, if agent **i** gets at least one of its goals achieved
- An agent has no preference over goals – it wants at least one to be achieved

Qualitative Coalitional Games

- A goal set **G'** is said to be **feasible** for coalition **C** if **G'** is one of the choices available to **C**, that is, if **G'** $\in$ **V(C)**
- A coalition is **successful** if **C** can cooperate in such a way that all of its members are satisfied – **C** is successful if there is some feasible goal set **G'** for **C** such that **G'** satisfies every member of **C**

Function **V** Representation

- Based on propositional logic
- System in which goal g_1 can be achieved by a coalition C iff C contains agent 1 and either 2 or 3 or both
- This characteristic function is represented using propositional logic:
- $g_1 \rightarrow (1 \wedge (2 \vee 3))$

Function **V** Representation

- Define a propositional formula, in which Boolean variables in the formula correspond to agents and goals
- **G’** $\in$ **V(C)** iff the formula ϕ, which represent **V**, evaluates to true under the valuation that makes the boolean variables corresponding to **G’** and **C** **true**, and all other variables **false**

Veto Player

- If **i** is a **veto** player for **j**, then we can say that **j** is dependent on **i**
- **Mutual dependence** – every player in a coalition is a veto player for every other player

Coalition Structure Formation

- A **coalition structure** is a partition of the overall set of agents **Ag** into mutually disjoint coalitions

Example

- 3 agents Ag = {1, 2, 3}
- 7 possible coalitions
- {1}, {2}, {3}, {1, 2}, {1, 3}, {2, 3}, {1, 2, 3}
- 5 possible coalition structures
- {{1}, {2}, {3}}
- {{1}, {2, 3}}
- {{2}, {1, 3}}
- {{3}, {1, 2}}
- {{1, 2, 3}}

Optimal Coalition Structure

- Given a coalitional game **(Ag,v)**, we say that the **socially optimal coalition structure CS*** with respect to **(Ag,v)** is given by:
- **CS* = arg** $\mathbf{max}_{CS \in partitions\ of\ Ag}\ \boldsymbol{V(CS)}$, where
- **V(CS) =** $\sum_{C \in CS} \boldsymbol{v(C)}$

Observation

- There will be exponentially many coalitions structures in the number of agents
- Searching through the space of all possible coalition structures to find the optimal coalition structure is undesirable, and infeasible in the worst case

Example

- 4 agents Ag = {1, 2, 3, 4}
- 15 possible coalition structures
- At level 1, we have all the possible coalition structures that contain exactly one coalition
- There is just one possible coalition structure containing a single coalition:
- {{1, 2, 3, 4}}

Example

- At level 2, we have all possible coalition structures that contain exactly 2 coalitions – all possible ways of partitioning the set of agents {1, 2, 3, 4} into 2 disjoint sets
- At level 3, we have all possible coalition structures that contain exactly 3 coalitions
- At level 4, we have all possible coalition structures that contain exactly 4 coalitions

Example

- Suppose we restrict the search for a possible coalition structure to the levels 1 and 2
- Let CS' be the best coalition structure that we find in these 2 levels, and let CS* be the best coalition structure overall
- Let C* be the coalition with the highest value of all possible coalitions:
- C* = arg $\max_{C \subseteq Ag} v(C)$

Example

- The value of the best coalition structure we find in the first two levels of the coalition structure graph must be at least as much as the value of the best possible coalition: $V(CS') \geq v(C^*)$
- This is because every possible coalition appears in at least one coalition structure in the first two levels of the graph
- Assume the worst case: $V(CS') = v(C^*)$

Example

- Since V(CS') is the highest possible value of any coalition structure, and there are only n = 4 agents, then the highest possible value of V(CS*) would be:
- n * v(C*) = n * V(CS')
- In the worst possible case, the value of the best coalition structure would be at worst 1/n the value of the best, where n is the number of agents

Example

- Searching the first 2 levels of the graph does not guarantee to give the optimal coalition structure, it guarantees to give one that is no worse than 1/n of the optimal

Algorithm

- 1) Search the bottom two levels of the coalition structure graph, keeping track of the best coalition structure seen so far
- 2) Continue with a breadth-first search, starting at the top of the graph, keeping track of the best coalition structure seen so far, and continue until either there is no remaining time, or else all the graph not studied in step (1) was considered
- 3) Return the coalition structure with the highest value