

Cooperative/coalitional game theory

Vincent Conitzer

Cooperative/coalitional game theory

- There is a set of agents N
- Each subset (or **coalition**) S of agents can work together in various ways, leading to various utilities for the agents
- Cooperative/coalitional game theory studies which outcome will/should materialize
- Key criteria:
 - **Stability**: No coalition of agents should want to deviate from the solution and go their own way
 - **Fairness**: Agents should be rewarded for what they contribute to the group
- (“Cooperative game theory” is the standard name (distinguishing it from **noncooperative game theory**, which is what we have studied so far). However this is somewhat of a misnomer because agents still pursue their own interests. Hence some people prefer “coalitional game theory.”)

Example

- Three agents $\{1, 2, 3\}$ can go out for Indian, Chinese, or Japanese food
- $u_1(I) = u_2(C) = u_3(J) = 4$
- $u_1(C) = u_2(J) = u_3(I) = 2$
- $u_1(J) = u_2(I) = u_3(C) = 0$
- Each agent gets an additional unit of utility for each other agent that joins her
- Exception: going out alone always gives a total utility of 0
- If all agents go for Indian together, they get utilities $(6, 2, 4)$
- All going to Chinese gives $(4, 6, 2)$, all going to Japanese gives $(2, 4, 6)$
- Hence, the **utility possibility set** for $\{1, 2, 3\}$ is $\{(6, 2, 4), (4, 6, 2), (2, 4, 6)\}$
- For the coalition $\{1, 2\}$, the utility possibility set is $\{(5, 1), (3, 5), (1, 3)\}$ (why?)

Stability & the core

- $u_1(I) = u_2(C) = u_3(J) = 4$
- $u_1(C) = u_2(J) = u_3(I) = 2$
- $u_1(J) = u_2(I) = u_3(C) = 0$
- $V(\{1, 2, 3\}) = \{(6, 2, 4), (4, 6, 2), (2, 4, 6)\}$
- $V(\{1, 2\}) = \{(5, 1), (3, 5), (1, 3)\}$
- Suppose the agents decide to all go for Japanese together, so they get $(2, 4, 6)$
- 1 and 2 would both prefer to break off and get Chinese together for $(3, 5)$ – we say $(2, 4, 6)$ is **blocked** by $\{1, 2\}$
 - Blocking only occurs if there is a way of breaking off that would make **all** members of the blocking coalition happier
- The **core** [Gillies 53] is the set of all outcomes (for the **grand coalition** N of all agents) that are blocked by no coalition
- In this example, the core is **empty** (why?)
- In a sense, there is no stable outcome

Transferable utility

- Now suppose that utility is **transferable**: you can give some of your utility to another agent in your coalition (e.g., by making a payment)
- Then, all that we need to specify is a **value** for each coalition, which is the maximum total utility for the coalition
 - Value function also known as **characteristic function**
- Any vector of utilities that sums to the value is possible
- Outcome is in the core if and only if: every coalition receives a total utility that is at least its value
 - For every coalition C , $v(C) \leq \sum_{i \in C} u(i)$
- In above example,
 - $v(\{1, 2, 3\}) = 12$,
 - $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = 8$,
 - $v(\{1\}) = v(\{2\}) = v(\{3\}) = 0$
- Now the outcome (4, 4, 4) is possible; it is also in the core (why?) and in fact the unique outcome in the core (why?)

Emptiness & multiplicity

- Let us modify the above example so that agents receive no utility from being together (except being alone still gives 0)
 - $v(\{1, 2, 3\}) = 6$,
 - $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = 6$,
 - $v(\{1\}) = v(\{2\}) = v(\{3\}) = 0$
- Now the core is empty!
- Conversely, suppose agents receive 2 units of utility for each other agent that joins
 - $v(\{1, 2, 3\}) = 18$,
 - $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = 10$,
 - $v(\{1\}) = v(\{2\}) = v(\{3\}) = 0$
- Now lots of outcomes are in the core – $(6, 6, 6)$, $(5, 5, 8)$, ...
- When is the core guaranteed to be nonempty?
- What about uniqueness?

Superadditivity

- v is **superadditive** if for all coalitions A, B with $A \cap B = \emptyset$, $v(A \cup B) \geq v(A) + v(B)$
- Informally, the union of two coalitions can always act as if they were separate, so should be able to get at least what they would get if they were separate
- Usually makes sense
- Previous examples were all superadditive
- Given this, always efficient for grand coalition to form

Convexity

- A game is **convex** if for all coalitions A, B , $v(A \cup B) - v(B) \geq v(A) - v(A \cap B)$ (i.e., v is **supermodular**)
- One interpretation: the **marginal contribution** of an agent is increasing in the size of the set that it is added to
- Previous examples were not convex (why?)
- In convex games, core is always nonempty
- One easy-to-compute solution in the core: agent i gets $u(i) = v(\{1, 2, \dots, i\}) - v(\{1, 2, \dots, i-1\})$
 - **Marginal contribution** scheme
 - Works for **any** ordering of the agents

The Shapley value [Shapley 1953]

- The marginal contribution scheme is unfair because it depends on the ordering of the agents
- One way to make it fair: average over **all** possible orderings
- Let $MC(i, \pi)$ be the marginal contribution of i in ordering π
- Then i 's **Shapley value** is $\sum_{\pi} MC(i, \pi)/(n!)$
- Always in the core for convex games
- ... but not in general, even when core is nonempty, e.g.
 - $v(\{1, 2, 3\}) = v(\{1, 2\}) = v(\{1, 3\}) = 1$,
 - $v = 0$ everywhere else

Axiomatic characterization of the Shapley value

- The Shapley value is the unique solution concept that satisfies:
 - **Efficiency**: the total utility is the value of the grand coalition,
 $\sum_{i \in N} u(i) = v(N)$
 - **Symmetry**: two symmetric players must receive the same utility
 - **Dummy**: if $v(S \cup \{i\}) = v(S)$ for all S , then i must get 0
 - **Additivity**: if we add two games defined by v and w by letting $(v+w)(S) = v(S) + w(S)$, then the utility for an agent in $v+w$ should be the sum of her utilities in v and w
 - most controversial axiom

Computing a solution in the core

- Can use linear programming:
 - Variables: $u(i)$
 - Distribution constraint: $\sum_{i \in N} u(i) = v(N)$
 - Non-blocking constraints: for every S , $\sum_{i \in S} u(i) \geq v(S)$
- Problem: number of constraints exponential in number of players
- ... but if the input explicitly specifies the value of every coalition, polynomial in input size
- ... but is this practical?

A concise representation based on synergies

[Conitzer & Sandholm AIJ06]

- Assume superadditivity
- Say that a coalition S is **synergetic** if there do not exist A, B with $A \neq \emptyset, B \neq \emptyset, A \cap B = \emptyset, A \cup B = S, v(S) = v(A) + v(B)$
- Value of non-synergetic coalitions can be derived from values of smaller coalitions
- So, only specify values for synergetic coalitions in the input

A useful lemma

- Lemma: For a given outcome, if there is a blocking coalition S (i.e., $\sum_{i \in S} u(i) < v(S)$), then there is also a synergetic blocking coalition
- Proof:
 - WLOG, suppose S is the smallest blocking coalition
 - Suppose S is not synergetic
 - So, there exist A, B with $A \neq \emptyset, B \neq \emptyset, A \cap B = \emptyset, A \cup B = S, v(S) = v(A) + v(B)$
 - $\sum_{i \in A} u(i) + \sum_{i \in B} u(i) = \sum_{i \in S} u(i) < v(S) = v(A) + v(B)$
 - Hence either $\sum_{i \in A} u(i) < v(A)$ or $\sum_{i \in B} u(i) < v(B)$
 - I.e., either A or B must be blocking
 - Contradiction!

Computing a solution in the core under synergy representation

- Can again use linear programming:
 - Variables: $u(i)$
 - Distribution constraint: $\sum_{i \in N} u(i) = v(N)$
 - Non-blocking constraints: for every **synergetic** S , $\sum_{i \in S} u(i) \geq v(S)$
- Still requires us to know $v(N)$
- If we do not know this, computing a solution in the core is NP-hard
- This is because computing $v(N)$ is NP-hard
- So, the hard part is not the strategic constraints, but computing what the grand coalition can do
- If the game is **convex**, then a solution in the core can be constructed in polynomial time even without knowing $v(N)$

Other concise representations of coalitional games

- [Deng & Papadimitriou 94]: agents are vertices of a graph, edges have weights, value of coalition = sum of weights of edges in coalition
- [Conitzer & Sandholm 04]: represent game as sum of smaller games (each of which involves only a few agents)
- [Jeong & Shoham 05]: multiple rules of the form (1 and 3 and (not 4) \rightarrow 7), value of coalition = sum of values of rules that apply to it
 - E.g., the above rule applies to coalition {1, 2, 3} (so it gets 7 from this rule), but not to {1, 3, 4} or {1, 2, 5} (so they get nothing from this rule)
 - Generalizes the above two representations (but not synergy-based representation)

Nucleolus [Schmeidler 1969]

- Always gives a solution in the core if there exists one
- Always uniquely determined
- A coalition's **excess** $e(S)$ is $v(S) - \sum_{i \in S} u(i)$
- For a given outcome, list all coalitions' excesses in decreasing order
- E.g., consider
 - $v(\{1, 2, 3\}) = 6$,
 - $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = 6$,
 - $v(\{1\}) = v(\{2\}) = v(\{3\}) = 0$
- For outcome $(2, 2, 2)$, the list of excesses is 2, 2, 2, 0, -2, -2, -2 (coalitions of size 2, 3, 1, respectively)
- For outcome $(3, 3, 0)$, the list of excesses is 3, 3, 0, 0, 0, -3, -3 (coalitions $\{1, 3\}$, $\{2, 3\}$; $\{1, 2\}$, $\{1, 2, 3\}$, $\{3\}$; $\{1\}$, $\{2\}$)
- Nucleolus is the (unique) outcome that **lexicographically minimizes** the list of excesses
 - Lexicographic minimization = minimize the first entry first, then (fixing the first entry) minimize the second one, etc.

Marriage contract problem

[Babylonian Talmud, 0-500AD]

- A man has three wives
- Their marriage contracts specify that they should, respectively, receive 100, 200, and 300 in case of his death
- ... but there may not be that much money to go around...
- Talmud recommends:
 - If 100 is available, each agent (wife) gets $33 \frac{1}{3}$
 - If 200 is available, agent 1 gets 50, other two get 75 each
 - If 300 is available, agent 1 gets 50, agent 2 gets 100, agent 3 gets 150
- ?
- Define $v(S) = \max\{0, \text{money available} - \sum_{i \in N-S} \text{claim}(i)\}$
 - Any coalition can walk away and obtain 0
 - Any coalition can pay off agents outside the coalition and divide the remainder
- Talmud recommends the nucleolus! [Aumann & Maschler 85]