Lecture 1
Introduction and Definition of TU games

1.1 Introduction

The study of the theory of games gained popularity after the publication of the classical book "Theory of games and economic behaviour" by von Neumann and Morgenstern in 1944[6]. A game is a mathematical model of an interaction between rational agents. The model first contains the strategy space of any participants (i.e., the actions that are available). From these different possible action, one can build the set of all possible states of the interaction. Then, the model contains the rules of interaction. These rules describe what happens when participants take some valid action, i.e., they determine what state is reached when the participants take some actions. Each participant has her own preferences between the different states, and each participant acts so as to obtain the best possible outcome (this is what we call rational, and we will come back to this later). The field of strategic game studies conflicting interactions and has been made popular with the contributions of Nash. Cooperative games is another field that analyses *cooperation* between agents.

The basic framework of cooperative game was introduced by von Neumann and Morgenstern in [6] and is called the characteristic function. A *coalition* is simply a set of agents that interacts. The so called characteristic function provides a payoff to each coalition. It is important to note that the payoff is given to a coalition, not to individual agent. The natural question that arises is the following: if all the agents want to cooperate, how should they share the payoff? Actually, if the agents also need to decide whether they want to cooperate and with which other agents, we come up with the following two key questions:

- the *selection* problem: which coalitions are going to form?
- the *sharing* problem: once the members have formed a coalition (i.e., they have self-organized, interacted, and the coalition formed received a payoff), the problem is then how to distribute it to the different members of the coalition.

The goal of the course is to answer these two questions.

In some parts of the course, we will focus only on the sharing problem and we will assume that all the participants intend to cooperate, forming one coalition containing all the participants (we will call this coalition the grand coalition).

The solutions to the sharing problem are called solution concepts and they are based on different interpretations of fairness. Unfortunately, there is no unique and accepted solution concept. For example, one possible criterion is stability: participants should not have an incentive to “change” (we will make this notion more formal later, but for now, let us consider that no participant want to change of coalitions or to ask for a different share of the payoff). As we will see, not all games can satisfy the most natural criterion modeling stability. Consequently, many other solution concepts have been proposed. A large part of the course will be about introducing the different solution concepts and study their properties.

We will also study some interesting special classes of games (i.e., restrictions of the model). A class of game may be interested because of its applications. For example, the term coalition is often used in political science: parties may form alliances to obtain more power. Consequently, we will study a class of games that models voting situation. A classes of game can also be interested because of some special properties (e.g., a solution with some properties is guaranteed to exist). The properties can also be computational. However, with cooperative games, one needs to be careful when dealing with computational issues as the input of the game is by nature exponential. Indeed, one needs to reason about all possible coalitions, i.e., all possible subset of the set of agents. Consequently, there are some interesting issues in representing the games and computing a solution. There are also some interesting issues to use, in practice, some solution concepts.

Finally, we will study different variations about modeling of cooperation. So far, we have talked about sharing the value of a coalition. But in some cases, there is not really a value to share: by forming a coalition, the members are in a specific state of the world and experience a corresponding satisfaction (e.g. think about which group of people to talk with at a party).

The course will mainly focus on the game theoretic aspect of cooperative games, and we will also study AI related issues towards the end of the course. Here is a rough outline of the course.

Solution concepts	The core Games with coalition structure and the bargaining set The nucleolus The kernel The Shapley value
Specific class of games	Voting games Representation and complexity
different model	NTU games and hedonic games
misc	Coalition formation and related issues

There is no official textbook for this course. There will be provide a lecture note for each class as this one. Here are some sources I used to prepare the class:

- The last three chapters of book "A course in game theory" by Osborne and Rubinstein [3] are devoted to cooperative games. I will use some of this material for the lectures on the core, the bargaining set, the kernel, the nucleolus and the Shapley value.

- The book "An introduction to the theory of cooperative games"by Peleg and Sudhölter [4] contains a rigorous and precise treatment of cooperative games. I used this book for some precision, but it is a more advanced textbook.

- For computational aspects and some advanced topics, you can read "Computational Aspects of Cooperative Game Theory" [1] by Chalkiadakis, Elkind and Wooldridge.

- Whenever appropriate, I will also refer to article from the literature.

1.2 TU games

The game theory community has extensively studied the coalition formation problem [2, 3]. 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.

1.2.1 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.

Notations

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 $\omega(n^{\frac{n}{2}})$ [5].

TU games

1.2.1. DEFINITION. 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.

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.

Standard Assumptions: It is usually assumed that the value of the empty coalition $\emptyset$ is 0 (i,.e. $v(\emptyset) = 0$; we will make that assumption throughout the class. Moreover, it is often the case that the value of each coalition is non-negative (when agents make profits) or else that the value of each coalition is non-posititive (when the members share some costs). During this class, we will assume that for each coalition $\mathcal{C}$, $v(\mathcal{C}) \geq 0$. However, most of the definitions and results can be easily adapted.

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.,

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

Some types of TU games

We now describes some types of valuation functions. First, we introduce a notion that will be useful on many occasion: the notion of marginal contribution. It represent the contribution of an agent when it joins a coalition.

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

The maximal marginal contribution $mc_i^{max} = \max_{\mathcal{C} \subseteq N \setminus \{i\}} mc_i(\mathcal{C})$ can been seen as a threat that an agent can use against a coalition: the agent can threatens to leave its current coalition to join the coalition that produces mc_i^{max}, arguing that it is able to generate mc_i^{max} utils. The minimal marginal contribution $mc_i^{min} = \min_{\mathcal{C} \subseteq N \setminus \{i\}} mc_i(\mathcal{C})$

is a minimum acceptable payoff: if the agent joins any coalition, the coalition will benefit by at most mc_i^{min}, hence agent i should get at least this amount.

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.

Monotone: $\forall \mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq N, v(\mathcal{C}_1) \leq v(\mathcal{C}_2)$. For example, the valuation function of the majority game is monotone: when more agents join a coalition, they cannot turn the coalition from a winning to a losing one. Many games are monotone, however, we can imagine non-monotone games. For instance, the overhead caused by costs of communication or the effort to cooperate may be such that adding another agent may decrease the value of a coalition. Another example features two agents that dislike each other: the productivity of a coalition may decrease when both of them are members of the coalition.

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. Many games are super-additive. As we have assumed that the value of a coalition is positive, superadditivity implies monotonicity (but the converse is not necessarily true). In such games, social welfare is maximised by forming the grand coalition. Consequently, the agents have incentives to form 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: A valuation is *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.

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.

Solutions

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.

It will be useful to talk about the payoff obtained by the members of a coalition and we will use the notation $x(\mathcal{C}) = \sum_{i \in \mathcal{C}} x(i)$.

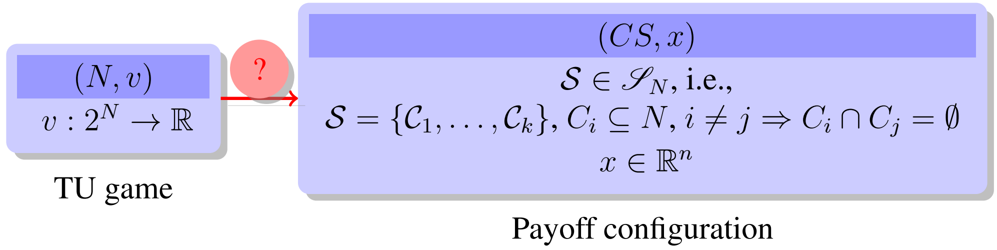

Figure 1.1: What is solving TU games?

We can finally formalize the solution of a TU game (N, v) by introducing the concept of *payoff configuration (PC)*. 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. The CS answers the *selection* problem when the payoff distribution answers the *sharing* problem. This is illustrated in Figure 1.1.

Let us now illustrate all these concepts with the following three-player TU game described in Table 1.1. 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: $\{\{1\}, \{2\}, \{3\}\}$, $\{\{1, 2\}, \{3\}\}$, $\{\{1\}, \{2, 3\}\}$, $\{\{2\}, \{1, 3\}\}$, $\{\{1, 2, 3\}\}$. This game is monotone and superadditive, but it is not convex.

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

Table 1.1: An example of a TU game

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 (the game is superadditive). 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? In this course, you will learn different ways to solve this problem. Unfortunately, as for many other games, we will see that there is not one unique best solution.

1.2.2 Rationality concepts

In this section, we discuss some desirable properties that link the coalition values to the agents' individual payoff. In other words, these properties are constraints that one would like to satisfy.

Feasible solution: First, one should not distribute more utility than is available. A payoff x is feasible when $\sum_{i \in N} x_i \leq v(N)$.

Anonymity: A solution is independent of the names of the agents. This is a pretty mild solution that will always be satisfied.

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. This is particularly relevant for superadditive games.

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 $\forall j \neq i,\ y_j \geq x_j\}$.

Reasonable from above: an agent should get at most its maximal threat, i.e., $x_i \leq mc_i^{max}$.

Reasonable from below: the agent should get at least its minimum acceptable reward $x_i \geq mc_i^{min}$.

Some more 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.

1.2.3. DEFINITION. 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.

1.2.4. DEFINITION. 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.

Bibliography

[1] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. *Computational aspects of cooperative game theory*. Morgan & Claypool, 2011.

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

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

[4] Bezalel Peleg and Peter Sudhölter. *Introduction to the theory of cooperative cooperative games*. Springer, 2nd edition, 2007.

[5] 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.

[6] John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 2^{nd} edition, 1947.