Computational Social Choice: Spring 2015

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Fair Allocation

Consider a set of agents and a set of goods. Each agent has their own preferences regarding the allocation of goods to agents to be selected.

- What constitutes a good allocation and how do we find it?

What goods? One or several goods? Available in single or multiple units? Divisible or indivisible? Can goods be shared? Are they static or do they change properties (e.g., consumable or perishable goods)?

What preferences? Ordinal or cardinal preference structures? Are monetary side payments possible, and how do they affect preferences? How are the preferences represented in the problem input?

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.

Plan for Today

This will be an introduction of fair allocation problems, focussing on multiple *indivisible goods* (also single-unit, non-sharable, static) for which agents express their preferences in terms of utility functions:

- Measuring *fairness* (and efficiency) of allocations
- Basic *complexity* results
- Allocation by means of *negotiation*

Most of this material is covered in my lecture notes cited below.

Recall that we've already talked about *cake cutting* (*divisible goods*).

U. Endriss. *Lecture Notes on Fair Division*. Institute for Logic, Language and Computation, University of Amsterdam, 2009/2010.

What is a Good Allocation?

We start with a (partial) overview of criteria that have been proposed for deciding what makes a "good" allocation:

- Of course, there are application-specific criteria, e.g.:
  - "*the allocation allows the agents to solve the problem*"
  - "*the auctioneer has generated sufficient revenue*"

  Here we are interested in general criteria that can be defined in terms of the individual agent preferences (*preference aggregation*).

- As we shall see, such criteria can be roughly divided into *fairness* and (economic) *efficiency* criteria.

Notation and Terminology

- Let $\mathcal{N} = \{1, \ldots, n\}$ be a set of *agents* (or *players*, or *individuals*) who need to share several *goods* (or *resources*, *items*, *objects*).
- An *allocation* A is a mapping of agents to *bundles* of goods.
- Each agent $i \in \mathcal{N}$ has a *utility function* u_i, mapping allocations to the reals, to model their preferences.
  - Typically, u_i is first defined on bundles, so: $u_i(A) = u_i(A(i))$.
  - Discussion: preference intensity, interpersonal comparison
- An allocation A gives rise to a *utility vector* $\langle u_1(A), \ldots, u_n(A) \rangle$.

Pareto Efficiency

Agreement A is *Pareto dominated* by agreement A' if $u_i(A) \leqslant u_i(A')$ for all agents $i \in \mathcal{N}$ and this inequality is strict in at least one case.

An agreement A is *Pareto efficient* if there is no other feasible agreement A' such that A is Pareto dominated by A'.

The idea goes back to Vilfredo Pareto (Italian economist, 1848–1923).

Collective Utility Functions

A *collective utility function* (CUF) is a function $\mathrm{SW} : \mathbb{R}^n \to \mathbb{R}$ mapping utility vectors to the reals ("social welfare"). Examples:

- The *utilitarian* CUF measures the sum of utilities:

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \sum_{i \in \mathcal{N}} u_i(A)$$

- The *egalitarian* CUF reflects the welfare of the agent worst off:

$$\mathrm{SW}_{\mathrm{egal}}(A) \quad = \quad min\{u_i(A) \mid i \in \mathcal{N}\}$$

- The *Nash* CUF is defined via the product of individual utilities:

$$\mathrm{SW}_{\mathrm{nash}}(A) \quad = \quad \prod_{i \in \mathcal{N}} u_i(A)$$

Remark: The Nash (like the utilitarian) CUF favours increases in overall utility, but also inequality-reducing redistributions ($2 \cdot 6 < 4 \cdot 4$).

Envy-Freeness

An allocation is called *envy-free* if no agent would rather have one of the bundles allocated to any of the other agents:

$$u_i(A(i)) \quad \geqslant \quad u_i(A(j))$$

Recall that $A(i)$ is the bundle allocated to agent i in allocation A.

<u>Remark:</u> Envy-free allocations do not always *exist* (at least not if we require either complete or Pareto efficient allocations).

Allocation of Indivisible Goods

We refine our formal framework as follows:

- Set of *agents* $\mathcal{N} = \{1, \ldots, n\}$ and finite set of indivisible *goods* $\mathcal{G}$.
- An *allocation* A is a partitioning of $\mathcal{G}$ amongst the agents in $\mathcal{N}$.
  Example: $A(i) = \{a, b\}$ — agent i owns items a and b
- Each agent $i \in \mathcal{N}$ has got a *utility function* $u_i : 2^{\mathcal{G}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy

How can we find a socially optimal allocation of goods?

Welfare Optimisation

How hard is it to find an allocation with maximal social welfare?

Rephrase this *optimisation problem* as a *decision problem:*

WELFARE OPTIMISATION (WO)
Instance: $\langle \mathcal{N}, \mathcal{G}, \mathcal{U} \rangle$ and $K \in \mathbb{Q}$
Question: Is there an allocation A such that $\mathrm{SW}_{\mathrm{util}}(A) > K$?

Unfortunately, the problem is intractable:

Theorem 1 WELFARE OPTIMISATION *is NP-complete, even when every agent assign nonzero utility to just a single bundle.*

Proof: NP-membership: we can check in polytime whether a given allocation A really has social welfare $> K$. NP-hardness: next slide. ✓

This seems to have first been stated by Rothkopf et al. (1998).

M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 44(8):1131–1147, 1998.

Proof of NP-hardness

By reduction to SET PACKING (known to be NP-complete):

> SET PACKING
> **Instance:** Collection $\mathcal{C}$ of finite sets and $K \in \mathbb{N}$
> **Question:** Is there a collection of disjoint sets $\mathcal{C}' \subseteq \mathcal{C}$ s.t. $|\mathcal{C}'| > K$?

Given an instance $\mathcal{C}$ of SET PACKING, consider this allocation problem:

- Goods: each item in one of the sets in $\mathcal{C}$ is a good
- Agents: one for each set in $\mathcal{C}$ + one other agent (called agent 0)
- Utilities: $u_C(S) = 1$ if $S = C$ and $u_C(S) = 0$ otherwise;
  $u_0(S) = 0$ for all bundles S

That is, every agent values "its" bundle at 1 and every other bundle at 0.
Agent 0 values all bundles at 0.

Then every set packing corresponds to an allocation (with SW $= |\mathcal{C}'|$).
Vice versa, for every allocation there is one with the same SW corresponding to a set packing (give anything owned by agents with utility 0 to agent 0). ✓

Welfare Optimisation under Additive Preferences

Sometimes we can reduce complexity by restricting attention to problems with certain types of preferences.

A utility function $u : 2^{\mathcal{G}} \to \mathbb{R}$ is called *additive* if for all $G \subseteq \mathcal{G}$:

$$u(S) \quad = \quad \sum_{g \in S} u(\{g\})$$

The following result is almost immediate:

Proposition 2 WELFARE OPTIMISATION *is in P in case all individual preferences are additive.*

Proof: To compute an allocation with maximal social welfare, simply give each item to (one of) the agent(s) who value it the most. ✓

This works, because we have $\sum_i \sum_g u_i(\{g\}) = \sum_g \sum_i u_i(\{g\})$.
So the same restriction does not help for, say, the egalitarian or Nash CUF.

Aside: Preference Representation

So far we have focussed on very simplistic preferences ...

Example: Allocating 10 goods to 5 agents means $5^{10} = 9765625$ allocations and $2^{10} = 1024$ bundles for each agent to think about.

So we need to choose a good *language* to compactly represent preferences over such large numbers of alternative bundles, e.g.:

- Logic-based languages (weighted goals)
- Bidding languages for combinatorial auctions (OR/XOR)
- Program-based preference representation (straight-line programs)
- CP-nets and CI-nets (for ordinal preferences)

The choice of language affects both *algorithm design* and *complexity*.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Distributed Approach

Instead of devising algorithms for computing a socially optimal allocation in a centralised manner, we now want agents to be able to do this in a distributed manner by contracting deals locally.

- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may come with a number of *side payments* to compensate some of the agents for a loss in utility. A *payment function* is a function $p : \mathcal{N} \to \mathbb{R}$ with $p(1) + \cdots + p(n) = 0$.

  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* €5, while agent j *receives* €5.

Negotiating Socially Optimal Allocations

We are not going to talk about designing a concrete negotiation protocol, but rather study the framework from an abstract point of view. The main question concerns the relationship between

- the *local view:* what deals will agents make in response to their individual preferences?; and
- the *global view:* how will the overall allocation of goods evolve in terms of social welfare?

We will go through this for one set of assumptions regarding the local view and one choice of desiderata regarding the global view.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

The Local/Individual Perspective

A rational agent (who does not plan ahead) will only accept deals that improve her individual welfare:

- A deal $\delta = (A, A')$ is called *individually rational* (IR) if there *exists* a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{N}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.

That is, an agent will only accept a deal if it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility).

The Global/Social Perspective

Suppose that, as system designers, we are interested in maximising *utilitarian social welfare:*

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \sum_{i \in \mathcal{N}} u_i(A(i))$$

Observe that there is no need to include the agents' monetary balances into this definition, because they'd always add up to 0.

While the local perspective is driving the negotiation process, we use the global perspective to assess how well we are doing.

Example

Let $\mathcal{A} = \{ann, bob\}$ and $\mathcal{G} = \{chair, table\}$ and suppose our agents use the following utility functions:

$$\begin{array}{rclrcl} u_{ann}(\emptyset) & = & 0 & u_{bob}(\emptyset) & = & 0 \\ u_{ann}(\{chair\}) & = & 2 & u_{bob}(\{chair\}) & = & 3 \\ u_{ann}(\{table\}) & = & 3 & u_{bob}(\{table\}) & = & 3 \\ u_{ann}(\{chair, table\}) & = & 7 & u_{bob}(\{chair, table\}) & = & 8 \end{array}$$

Furthermore, suppose the initial allocation of goods is A_0 with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \emptyset$.

Social welfare for allocation A_0 is 7, but it could be 8. By moving only a *single* good from agent ann to agent bob, the former would lose more than the latter would gain (not individually rational).

The only possible deal would be to move the whole *set* $\{chair, table\}$.

Convergence

The good news:

Theorem 3 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Discussion: Agents can act *locally* and need not be aware of the global picture (convergence is guaranteed by the theorem).

Discussion: Other results show that (a) arbitrarily complex deals might be needed and (b) paths may be exponentially long. Still NP-hard!

T. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. Proc. AAAI Spring Symposium 1998.

So why does this work?

The key to the proof is the insight that IR deals are exactly those deals that increase social welfare:

- **Lemma 4** *A deal $\delta = (A, A')$ is individually rational if and only if $\mathrm{SW}_{\mathrm{util}}(A) < \mathrm{SW}_{\mathrm{util}}(A')$.*

  Proof: ($\Rightarrow$) Rationality means that overall utility gains outweigh overall payments (which are $= 0$).

  ($\Leftarrow$) The social surplus can be divided amongst all agents by using, say, the following payment function:

$$p(i) \;=\; u_i(A') \;-\; u_i(A) \;-\; \underbrace{\frac{\mathrm{SW}_{\mathrm{util}}(A') - \mathrm{SW}_{\mathrm{util}}(A)}{|\mathcal{N}|}}_{>0} \qquad \checkmark$$

Thus, as SW increases with every deal, negotiation must *terminate*. Upon termination, the final allocation A must be *optimal*, because if there were a better allocation A', the deal $\delta = (A, A')$ would be IR.

Summary

Fairness and *efficiency* criteria introduced:

- Pareto efficiency (very basic)
- Utilitarian, egalitarian, Nash collective utility
- Envy-freeness

We have seen that finding a fair/efficient allocation in case of indivisible goods gives rise to a combinatorial optimisation problem.

Two approaches:

- *Centralised:* Give a complete specification of the problem to an optimisation algorithm. Often *intractable*.

- *Distributed:* Try to get the agents to solve the problem.
  For certain fairness criteria and certain assumptions on agent behaviour, we can predict *convergence* to an optimal state.