

Computational Social Choice: Spring 2017

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

This is an introduction to *fair allocation* problems for *indivisible goods* for which agents express their preferences in terms of *utility functions*:

- measuring *fairness* (and efficiency) of allocations
- basic *complexity* results
- *protocols* to interactively determine a good allocation

Most of this material is also covered in my lecture notes cited below. For more, consult the *Handbook*.

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

U. Endriss. *Lecture Notes on Fair Division*. ILLC, University of Amsterdam, 2009.

S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair Allocation of Indivisible Goods. In F. Brandt et al. (eds.), *Handbook of Computational Social Choice*. CUP, 2016.

Notation and Terminology

Let $N = \{1, \dots, n\}$ be a group 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 $i \in N$ has a *utility function* u_i , mapping allocations to the reals, to model her preferences.

- Typically, u_i is first defined on bundles, so: $u_i(A) = u_i(A(i))$.
- Discussion: preference intensity, interpersonal comparison

Every allocation A gives rise to a *utility vector* $(u_1(A), \dots, u_n(A))$.

Exercise: *What would be a good allocation? Fair? Efficient?*

Collective Utility Functions

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

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

$$SW_{\text{util}}(A) = \sum_{i \in N} u_i(A)$$

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

$$SW_{\text{egal}}(A) = \min\{u_i(A) \mid i \in N\}$$

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

$$SW_{\text{nash}}(A) = \prod_{i \in N} u_i(A)$$

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

Pareto Efficiency

Some criteria require only *ordinal* comparisons ...

Allocation A is *Pareto dominated* by allocation A' if $u_i(A) \leq u_i(A')$ for all agents $i \in N$ and this inequality is strict in at least one case.

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

Envy-Freeness

An allocation is *envy-free* if no agent would want to swap her own bundle with the bundle assigned to one of the other agents:

$$u_i(A(i)) \geq u_i(A(j))$$

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

Exercise: *Show that for some scenarios there exists no allocation that is both envy-free and Pareto efficient.*

Allocation of Indivisible Goods

We refine our formal framework as follows:

- Set of *agents* $N = \{1, \dots, n\}$ and finite set of indivisible *goods* G .
- An *allocation* A is a partitioning of G amongst the agents in N .
Example: $A(i) = \{a, b\}$ — agent i owns items a and b
- Each agent $i \in N$ has got a *utility function* $u_i : 2^G \rightarrow \mathbb{R}$, giving rise to a profile of utility functions $\mathbf{u} = (u_1, \dots, u_n)$.
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?

- Could think of this as a *combinatorial optimisation* problem.
- Or devise a *protocol* to let agents solve the problem interactively.

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: $\langle N, G, \mathbf{u} \rangle$ and $K \in \mathbb{Q}$

Question: Is there an allocation A such that $\text{SW}_{\text{util}}(A) > K$?

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 “her” 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 $u : 2^G \rightarrow \mathbb{R}$ is called *additive* if for all $S \subseteq G$:

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

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. ✓

Remark: 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.

Negotiating Socially Optimal Allocations

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.

- We are given some *initial allocation* A_0 .
- 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 $p : N \rightarrow \mathbb{R}$ with $p(1) + \dots + p(n) = 0$.

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

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 $i \in N$, 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*:

$$SW_{\text{util}}(A) = \sum_{i \in N} u_i(A(i))$$

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.

Exercise: *How well (or how badly) do you expect this to work?*

Example

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

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

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 IR iff $SW_{\text{util}}(A) < SW_{\text{util}}(A')$.*

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:

$$p(i) = u_i(A') - u_i(A) - \underbrace{\frac{SW_{\text{util}}(A') - SW_{\text{util}}(A)}{|N|}}_{> 0} \quad \checkmark$$

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.

Related Work

Many ways in which this can be (and has been) taken further:

- other social objectives? / other local criteria?
- what types of deals needed for what utility functions?
- path length to convergence?
- other types of goods: sharable, nonstatic, ... ?
- negotiation on a social network?

For several combinations of the above there still are open problems.

Summary

We saw several *fairness* and *efficiency* criteria for allocation problems:

- utilitarian, egalitarian, Nash collective utility
- Pareto efficiency and envy-freeness

We saw that finding a good allocation in case of indivisible goods gives rise to a combinatorial optimisation problem. Two approaches:

- *Centralised*: Give a complete specification of the problem to a combinatorial optimisation algorithm. Often *intractable*.
- *Distributed*: Try to get the agents to solve the problem. For some parameters we can guarantee *convergence* to an optimum.

Many more interesting topics around fair allocation. Selection:

- compact preference representation (e.g., via *weighted goals*)
- other fairness criteria (such as the *maximin fare share*)
- axiomatic characterisation of criteria (e.g., via *scale invariance*)
- other protocols (such as *picking sequences*)