

Computational Social Choice: Spring 2019

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

We continue our discussion of the problem of finding a *fair allocation* of a number of *indivisible goods* to several agents.

Today we focus on algorithmic concerns:

- *complexity* of computing a socially optimal allocation in a *centralised manner*, given the agents' preferences
- finding socially optimal allocations by asking agents to follow a given *interaction protocol* (our example: *picking sequences*)
- asking agents to *negotiate* socially optimal allocations in a fully *distributed manner*, by contracting local deals

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.

The Model

We stick to our familiar model of fair allocation with indivisible goods:

Let $N = \{1, \dots, n\}$ be a set of *agents* and G a finite set of *goods*.

Every agent $i \in N$ has a *utility* (or *valuation*) *function* $u_i : 2^G \rightarrow \mathbb{R}$, indicating how much she values any given *bundle* of goods.

An *allocation* $A : N \rightarrow 2^G$ is a mapping from agents to bundles that respects $A(i) \cap A(j) = \emptyset$ for $i \neq j$ and $A(1) \cup \dots \cup A(n) = G$.

Every allocation A induces a *utility vector* $(u_1(A(1)), \dots, u_n(A(n)))$.

We want to choose an allocation based on this.

Maximising Social Welfare

Recall that the *utilitarian social welfare* on an allocation is the sum of the individual utilities it induces. *How hard is it to maximise USW?*

MAXIMUM UTILITARIAN SOCIAL WELFARE (MAXUSW)

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

Question: Is there an allocation A such that $\sum_{i \in N} u_i(A(i)) > K$?

Unfortunately, this problem is intractable:

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

Proof: *NP-membership:* can check in poly-time whether a proposed solution 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 from SET PACKING (which is known to be NP-hard):

SET PACKING

Input: 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 scenario:

- 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

Thus: 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 USW = $|\mathcal{C}'|$).

And for every allocation there is one with the same USW corresponding to a set packing (give anything owned by agents with utility 0 to agent 0). ✓

Special Case: Additive Utility Functions

Sometimes we can reduce the complexity by restricting attention to allocation scenarios with specific types of preferences.

Recall: $u_i : 2^G \rightarrow \mathbb{R}$ is *additive* if $u(S) = \sum_{x \in S} u_i(\{x\})$ for all $S \subseteq G$.

Now the generally intractable MAXUSW becomes tractable:

Proposition 2 MAXUSW can be decided in polynomial time in case all individual utility functions are *additive*.

Exercise: Can you see why?

Remark: This wouldn't work for, say, egalitarian or Nash social welfare.

More Complexity Results

To give you a rough idea of what other kinds of complexity results are out there, here are some representative examples (omitting precise statements regarding the representation of preferences employed):

- Finding an allocation with maximal *egalitarian* social welfare is NP-hard, even when all valuations are additive.
- Checking whether a given allocation is *Pareto efficient* is coNP-complete.
- Checking whether an *envy-free* allocation exists is NP-complete; checking whether an allocation that is both Pareto efficient and envy-free exists is even Σ_2^p -complete (NP with NP-oracle).

References to these results may be found in the MARA Survey.

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.

Picking Sequences

Sometimes we do not just want to treat fair allocation as a *centralised* combinatorial optimisation problem, but instead want agents to find good allocations themselves, following a suitable *interaction protocol*.

Example: Ask agents to take turns picking an item, until none are left. You may remember this from picking football teams at school.

Very low *complexity*. Very modest *communication* overheads. Great! But how *fair* is this? And why not, say, 123321 instead of 123123?

So let us analyse *picking sequences*:

- A *policy* $\pi : \{1, \dots, m\} \rightarrow N$, for $m = |G|$, fixes which agent is allowed to pick one of the remaining goods at any given time.
- Focus on *additive* u_i with $u_i(x) := u_i(\{x\}) > 0$ for all $x \in G$.

S. Bouveret and J. Lang. A General Elicitation-Free Protocol for Allocating Indivisible Goods. IJCAI-2011.

Strict Alternation Guarantees EF1

The *strict alternation* (or *round-robin*) policy π is the first thing to try:

$$\underbrace{123 \dots n \ 123 \dots n \ 123 \dots}_{m}$$

Recall: An allocation is *envy-free up to one good* (EF1) if no agent i is envious of any other agent, as long as i can remove one item from j 's bundle. Maximal NSW allocations are EF1 (but hard to compute).

Good news:

Proposition 3 *When all agents have additive utility functions, then the strict alternation protocol ensures envy-freeness up to one good.*

Proof: If $i < j$ (i precedes j in the sequence), then i will always pick before j and thus not envy her. If $j < i$, remove the item j picked first from her bundle: then things are *as if* i had always picked before j . ✓

Social Welfare

Providing guarantees regarding social welfare is much harder.

To have a chance, we simplify further (beyond assuming additivity):

Every agent $i \in N$ has a strict preference ranking \succ_i on G and her utility for a good $x \in G$ depends only on \succ_i . Options:

- *Borda Scoring*: $u_i(x) = 1 + \#\{y \in G \mid x \succ_i y\}$
- *Lexicographic Scoring*: $u_i(x) = 2^{\#\{y \in G \mid x \succ_i y\}}$
- Any other *scoring function* $g : \{1, \dots, m\} \rightarrow \mathbb{R}^+$ that is nondecreasing in $\#\{y \in G \mid x \succ_i y\}$.

We consider two scenarios:

- *Full Correlation*: All agents have the exact same ranking \succ .
- *Full Independence*: Agents draw rankings independently from the uniform probability distribution over the set of all rankings.

Want to know: *Which policy maximises (expected) social welfare?*

Full Correlation: Utilitarian Social Welfare

For certain combinations of assumptions our question becomes trivial:

Fact 4 For *fully correlated* preferences and any scoring function g , *every policy* leads to the *same utilitarian social welfare*.

Exercise: *Can you see why?*

Full Correlation: Egalitarian Social Welfare

Good news:

Proposition 5 (Bouveret and Lang, 2011) *For fully correlated preferences and lexicographic scoring the policies that maximise egalitarian social welfare are exactly those of this form (for $m \geq n$):*

$$\sigma(1) \sigma(2) \cdots \sigma(n-1) \underbrace{\sigma(n) \cdots \sigma(n)}_{m-n+1 \text{ times}} \quad \text{for some permutation } \sigma : N \rightarrow N$$

Exercise: *Can you see why?*

S. Bouveret and J. Lang. A General Elicitation-Free Protocol for Allocating Indivisible Goods. IJCAI-2011.

Full Correlation: Egalitarian Social Welfare

For Borda scoring, there is no one-size-fits-all solution, but still:

Proposition 6 (Bouveret and Lang, 2011) *For fully correlated preferences and Borda scoring a policy that maximises egalitarian social welfare can be computed in pseudopolynomial time.*

Note: pseudopolynomial = polynomial w.r.t. unary repres. of n and m

Proof: Omitted. We basically have to cluster the numbers $\{1, \dots, m\}$ into n sets, in a way that maximises the lowest sum across all classes. As the range of possible sums we might encounter is polynomial in m (at most $\frac{m(m+1)}{2}$), this can be done using dynamic programming.

S. Bouveret and J. Lang. A General Elicitation-Free Protocol for Allocating Indivisible Goods. IJCAI-2011.

Full Correlation: Egalitarian Social Welfare

But for *arbitrary scoring functions*, the problem becomes intractable:

Proposition 7 (Bouveret and Lang, 2011) *For fully correlated preferences, deciding whether there exists a policy that ensures a given level K of egalitarian social welfare is NP-complete (even for $n = 2$).*

Proof: *Membership in NP* follows from the fact that we can compute ESW for a given policy π (the certificate) in polynomial time. ✓

NP-hardness can be proved by reduction from PARTITION (next slide).

S. Bouveret and J. Lang. A General Elicitation-Free Protocol for Allocating Indivisible Goods. IJCAI-2011.

Proof of NP-hardness

By reduction from PARTITION (which is known to be NP-hard):

PARTITION

Input: $(w_1, \dots, w_m) \in \mathbb{N}^m$ with $w_1 \leq w_2 \leq \dots \leq w_m$

Question: Is there a set $S \subseteq \{1, \dots, m\}$ such that $\sum_{j \in S} w_j = \sum_{j \notin S} w_j$?

Given an instance (w_1, \dots, w_m) of PARTITION, build an allocation scenario for m goods and two agents with common ranking \succ :

$$u_i(x) = w_k \text{ with } k = 1 + \#\{y \in G \mid x \succ y\}$$

Then achieving ESW of K for $K := \frac{1}{2} \cdot (w_1 + \dots + w_m)$ is possible iff (w_1, \dots, w_m) is a YES-instance of PARTITION. Using this policy π :

$$\pi(j) = \begin{cases} 1 & (m - j + 1) \in S \\ 2 & \text{otherwise} \end{cases}$$

Thus, finding a suitable policy is at least as hard as PARTITION. ✓

Full Independence

If the agents do not all have the same utility function (full correlation), the analysis becomes more difficult, even assuming *full independence*.

But a few things are known for *Borda scoring* (details omitted):

- For $m \rightarrow \infty$ (number of goods), any *balanced* policy (number of picks any two agents had never differs by more than 1) maximises expected *utilitarian* and *egalitarian SW* (Bouveret and Lang, 2011).
- For *two agents*, the *strict alternation* policy maximises expected *utilitarian social welfare* (Kalinowski et al., 2013).
- Bouveret and Lang (2011) also compute optimal policies for certain small values of n and m by exhaustive search.

S. Bouveret and J. Lang. A General Elicitation-Free Protocol for Allocating Indivisible Goods. IJCAI-2011.

T. Kalinowski, N. Narodytska, and T. Walsh. A Social Welfare Optimal Sequential Allocation Procedure. IJCAI-2013.

Negotiating Socially Optimal Allocations

Rather than computing optimal allocations in a centralised manner or prescribing a fixed protocol, we now would like agents to settle on an allocation 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 however does not plan ahead) will only accept deals that improve her individual welfare:

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

Thus, 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 measure the quality of the allocations A the agents negotiate in terms of *utilitarian social welfare*:

$$\sum_{i \in N} u_i(A(i))$$

Remark: There is no need to include the agents' monetary balances into this definition, because those would 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

Good news:

Theorem 8 (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 9 *A deal $\delta = (A, A')$ is IR iff $USW(A) < USW(A')$.*

Proof: (\Rightarrow) Individual rationality of δ implies that globally *utility gains outweigh payments* (which are = 0). ✓

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

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

Thus, as USW 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

This concludes our review of *fair allocation* problems.

In terms of *modelling* such problems, we discussed these issues:

- goods: divisible (“cake”) vs. indivisible, money vs. no money
- preferences: utilities vs. orders, additivity, representation

In terms of *requirements*, we saw many different notions of fairness.

In terms of *solutions*, today we saw three kinds of approaches:

- computing socially optimal allocations in a centralised manner
- finding such allocation by asking agents to follow a protocol
- finding such allocations by letting agents negotiate

What next? We will switch to the social choice scenario of voting.