Computational Social Choice: Spring 2019

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

Plan for Today

Our first major topic for this course is the *fair allocation* of one or more goods to two or more agents.

Cake cutting was a first instance (with a single infinitely divisible item).

We now switch to *indivisible goods* and today focus on *axiomatics:*

- definition of the *formal model:* agents, goods, utility, allocations
- many examples for possible *social objectives:* what is *fair*?
- formal properties for a few: first instance of *axiomatic method*

Also: *nature of preferences* and *compact preference representation*

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

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.

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

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

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

- Remark: u_i need not (but could) be *additive* ($u_i(S) = \sum_{x \in S} u_i(\{x\})$).
- Discussion: This presupposes *preference intensity* makes sense.
  Ideally, we'd just use $\succcurlyeq_i$ (with $S \succcurlyeq_i S' \;\Leftrightarrow\; u_i(S) \geqslant u_i(S')$).
- Discussion: Also presupposes *interpersonal comparisons* make sense.

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

Every allocation A induces a *utility vector* $(u_1(A(1)), \ldots, u_n(A(n)))$.
We want to chose an allocation based on this.

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

Utilitarian Social Welfare

Given the utilities of the individual agents, we can define a notion of social welfare and aim for an agreement that maximises social welfare.

Classical utilitarianism (defended by Jeremy Bentham, 1748–1832) suggests to pick an allocation that maximises the sum of utilities:

$$\operatorname*{argmax}_{A \in \mathrm{Alloc}} \sum_{i \in N} u_i(A(i))$$

This sum is known as the *utilitarian social welfare* of allocation A.

Remark: Maximising USW amounts to maximising *average utility*.

Exercise: *Is maximising USW a good social objective?*

Egalitarian Social Welfare

An alternative approach is to pick an allocation that maximises the utility of the agent who is worst off:

$$\operatorname*{argmax}_{A \in \text{Alloc}} \min\{\, u_i(A(i)) \mid i \in N \,\}$$

This minimum is called the *egalitarian social welfare* of allocation A.

ESW is inspired by the philosophy of John Rawls (1921–2002) and was formally developed, amongst others, by Amartya Sen since the 1970s (Nobel Prize in Economic Sciences in 1998).

A refinement is the *leximin ordering:* first compute all allocations that maximise the lowest utility; then, from within this set, compute all allocations that maximise the utility of the next-poorest agent; repeat.

J. Rawls. *A Theory of Justice*. Oxford University Press, 1971.

A.K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.

The Nash Product

We tried adding utilities and computing minima. *Why not multiply?*

$$\operatorname*{argmax}_{A \in \text{Alloc}} \prod_{i \in N} u_i(A(i))$$

This product is called the *Nash product* (or *Nash social welfare*) of A. Named after John Nash (Nobel Prize in Economic Sciences in 1994).

Intuition: NSW favours increases in overall utility (just like USW), but also inequality-reducing redistributions of utility ($2 \cdot 6 < 4 \cdot 4$).

Remark: NSW only makes sense when utilities are positive.

Collective Utility Functions

USW, ESW, and NSW are examples for *collective utility functions*.

A CUF is a function $F : \mathbb{R}^n \to \mathbb{R}$ we can use to score utility vectors.

Remark: To apply a CUF to allocations, first compute utility vectors.

An important family of CUFs:

Given allocation A and $k \leqslant n$, let $u_k^\star(A)$ be the k-lowest number in the utility vector $\boldsymbol{u}(A) = (u_1(A(1)), \ldots, u_n(A(n)))$.

The *k-rank dictator CUF* maps any $\boldsymbol{u}(A)$ to its $u_k^\star(A)$.

Examples:

- For $k = 1$ we get the familiar *egalitarian CUF*.
- For $k = n$ we get the *elitist CUF*.
- For $k = \lfloor \frac{n+1}{2} \rfloor$ we get the *median-rank dictator CUF*.

The Axiomatic Method

So how do you choose a suitable social objective?

The central approach in SCT is the *axiomatic method:*

- identify and formalise normatively appealing properties ("*axioms*")
- check which methods satisfy your axioms

Let us focus on *social welfare orderings:*

> A SWO is a binary relation $\succcurlyeq$ on the set of all utility vectors that is reflexive, transitive, and complete.

Every CUF F *induces* a SWO: $\boldsymbol{u}(A) \succcurlyeq \boldsymbol{u}(A')$ if $F(\boldsymbol{u}(A)) \geqslant F(\boldsymbol{u}(A'))$. The opposite is true in our specific setting, but not in general.

Remark: Today we will be content with discussing some attractive axioms and checking which of our SWOs satisfy them. But stronger results, fully characterising certain SWOs in terms of axioms, exist.

The Pigou-Dalton Principle

A fair SWO should encourage inequality-reducing utility redistributions.

Axiom 1 (PD) *A SWO $\succcurlyeq$ respects the Pigou-Dalton principle if, for all $\boldsymbol{u}(A), \boldsymbol{u}(A') \in \mathbb{R}^n$, it is the case that $\boldsymbol{u}(A) \succcurlyeq \boldsymbol{u}(A')$ holds whenever there are two agents $i, j \in N$ such that:*

- $u_k(A) = u_k(A')$ *for all* $k \in N \setminus \{i, j\}$ — *only* i *and* j *are involved*
- $u_i(A) + u_j(A) \geqslant u_i(A') + u_j(A')$ — *at least mean-preserving*
- $|u_i(A) - u_j(A)| < |u_i(A') - u_j(A')|$ — *inequality-reducing*

The idea is due to Arthur C. Pigou (British economist, 1877–1959) and Hugh Dalton (British economist and politician, 1887–1962).

Example: The leximin ordering satisfies the Pigou-Dalton principle.

Zero Independence

If agents initially enjoy very different utilities, it may not be meaningful to use their absolute utilities to assess social welfare, but rather their *relative* gains or losses in utility. So a desirable property of a SWO may be to be independent of what individual agents consider "zero" utility.

Axiom 2 (ZI) *A SWO $\succcurlyeq$ is zero-independent if $\boldsymbol{u}(A) \succcurlyeq \boldsymbol{u}(A')$ entails $(\boldsymbol{u}(A) + \boldsymbol{w}) \succcurlyeq (\boldsymbol{u}(A') + \boldsymbol{w})$ for all $\boldsymbol{u}(A), \boldsymbol{u}(A'), \boldsymbol{w} \in \mathbb{R}^n$.*

Example: The utilitarian SWO is ZI, but the egalitarian SWO is not.

Remark: A SWO satisfies ZI iff it is the utilitarian SWO. Refer to the book by Moulin (1988) for a precise statement of this result.

H. Moulin. *Axioms of Cooperative Decision Making*. Econometric Society Monographs, Cambridge University Press, 1988.

Scale Independence

Different agents may measure their utility using different "currencies". So a desirable property of a SWO may be to be independent of this.

Assumption: All utilities involved are positive.

Notation: $\boldsymbol{v} \cdot \boldsymbol{w} = (v_1 \cdot w_1, \ldots, v_n \cdot w_n)$ for $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}^n$.

Axiom 3 (SI) *A SWO* $\succcurlyeq$ *is scale-independent if* $\boldsymbol{u}(A) \succcurlyeq \boldsymbol{u}(A')$ *entails* $(\boldsymbol{u}(A) \cdot \boldsymbol{w}) \succcurlyeq (\boldsymbol{u}(A') \cdot \boldsymbol{w})$ *for all* $\boldsymbol{u}(A), \boldsymbol{u}(A'), \boldsymbol{w} \in (\mathbb{R}^+)^n$.

Example: Neither the utilitarian nor the egalitarian SWO satisfy SI. However, the Nash SWO does.

Remark: By a similar result as the one mentioned before, a SWO satisfies SI iff it is the Nash SWO.

Independence of the Common Utility Pace

Another desirable property of a SWO may be that we would like to be able to make social welfare judgements without knowing what kind of "tax" members of society will have to pay.

Axiom 4 (ICP) *A SWO $\succcurlyeq$ is independent of the common utility pace if $\boldsymbol{u}(A) \succcurlyeq \boldsymbol{u}(A')$ entails $g(\boldsymbol{u}(A)) \succcurlyeq g(\boldsymbol{u}(A'))$ for all $\boldsymbol{u}(A), \boldsymbol{u}(A') \in \mathbb{R}^n$ and every increasing bijection $g : \mathbb{R} \to \mathbb{R}$ (applied component-wise).*

For a SWO satisfying ICP only relative interpersonal comparisons (such as $u_i(A) \leqslant u_j(A')$) matter, but not their (cardinal) intensities (such as $u_i(A) - u_j(A')$).

Example: The utilitarian SWO does *not* satisfy ICP. The egalitarian SWO does. Any k-rank dictator SWO does.

Pareto Efficiency

Allocation A is *Pareto dominated* by allocation A' if $u_i(A) \leqslant u_i(A')$ 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'.

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

Discussion:

- Pareto efficiency is very often considered a minimum requirement for any decent allocation. It is a very weak criterion.
- Only the ordinal component of utility is needed to check Pareto efficiency (no preference intensity, no interpersonal comparison).

Envy-Freeness

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

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

Remark: Also only requires the ordinal component of utility. Good.

Remark: Presupposes agents care not just about their own bundle.

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

Relaxing Envy-Freeness: Degrees of Envy

As we cannot always get envy-freeness, we may try to *minimise envy*.

Several definitions for the *degree of envy* of an allocation A:

- number of agents i experiencing some envy
- number of envy-relationships ((i, j) with $u_i(A(i)) < u_i(A(j)$)
- maximum number of peers envied by any individual
- sum of all envies experienced by anyone
- etc.

Refer to Chevaleyre et al. (2017) for a systematic approach.

Y. Chevaleyre, U. Endriss, and N. Maudet. Distributed Fair Allocation of Indivisible Goods. *Artificial Intelligence*, 2017.

Relaxing Envy-Freeness: EF1

An allocation A is called *envy-free up to one good* (EF1) if, whenever an agent i envies another agent j, then that envy can be eliminated by removing one item from j's bundle. So this must hold for all $i, j \in N$:

$$u_i(A(i)) \geqslant u_i(A(j) \setminus \{x\}) \text{ for some item } x \in G$$

Caragiannis et al. (2019) show that for scenarios with *additive* utility functions every allocation with maximal *Nash social welfare* is *EF1* (!).

I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, and J. Wang. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 2019 (in press).

The Maximin Share

Recall how for we had defined proportional fairness for cake cutting.

Even for additive utility functions, this definitions is not overly helpful when allocating indivisible goods. Exercise: *Can you see why?*

But here is a similar idea, inspired by the cut-and-choose protocol:

> Suppose you're asked to split G into n bundles, then everyone picks a bundle, with you going last (so might get the worst). The utility you can guarantee yourself is your *maximin share*.

So another social objective would be to find an allocation that gives everyone at least their maximin share. (Not always possible.)

E. Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 2011.

The Lorenz Curve

Ideally, every single agent will enjoy exactly the same level of utility. The Lorenz curve is a way to visualise how far we are from this ideal.

Let $\boldsymbol{u}^\star(A)$ be the *ordered utility vector* of allocation A. Then this is the total utility of the k poorest agents:

$$L_k(A) \quad = \quad \sum_{i=1}^{k} u_i^\star(A(i))$$

Remark: $L_1(A)$ is the ESW and $L_n(A)$ is the USW of A.

The vector $(L_1(A), \ldots, L_n(A))$ is called the *Lorenz curve* of A. Want an allocation with a Lorentz curve not dominated by any other. Ideal would be the *line of perfect equality* $(1 \cdot \frac{L_n(A)}{n}, \ldots, n \cdot \frac{L_n(A)}{n})$.

Exercise: *Which utility vector is better?* $(1, 2, 7, 7, 8)$ *or* $(1, 3, 5, 6, 10)$*?*

Inequality Indices

An *inequality index* is a function mapping allocations to $[0, 1]$, with 0 representing perfect equality and 1 representing complete inequality.

Two popular indices:

- *Gini index* $=$ *area* between line of perfect equality and Lorenz curve (divided by a suitable normalisation factor)
- *Robin Hood index* $=$ maximal *distance* between line of perfect equality and Lorenz curve (also normalised)

Now we can discern $(1, 2, 7, 7, 8)$ and $(1, 3, 5, 6, 10)$: the former is better according to Gini, the latter according to Robin Hood.

Preference Representation Languages

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 a good *language* to compactly represent preferences over such large numbers of alternative bundles. Some options:

- Logic-based languages (weighted goals)
- Bidding languages for combinatorial auctions (OR/XOR)
- CP-nets and CI-nets (for ordinal preferences)

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

Furthermore, different languages may have different *expressive power* or (arguably) enjoy different *cognitive relevance*.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 2008.

Example for a Language: Weighted Goals

Think of items in G as propositional variables. So: models $=$ bundles

Agents can now express *goals* as propositional formulas and attach numerical *weights* to them to signal importance. Examples:

$(p \wedge \neg q, 4)$ "I want item p but not item q" (medium important)
$(p \wedge q \rightarrow r, 7)$ "If I get p and q, I also want r" (very important)

Any *goalbase* Goals_i of such weighted goals defines a utility function u_i:

$$u_i(S) \quad = \sum_{(\varphi,w)\in\mathrm{Goals}_i} \mathbb{1}_{S\models\varphi} \cdot w$$

We get a *family of languages* $\mathcal{L}(H_1, H_2)$, one for every restriction H_1 on formulas and H_2 on weights: $\mathcal{L}(clauses, \mathbb{N})$, $\mathcal{L}(literals, \mathbb{R})$, ...

The results on the next two slides are taken from the paper cited below.

J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing Utility Functions via Weighted Goals. *Mathematical Logic Quarterly*, 2009.

Expressive Power

Proposition 1 *$\mathcal{L}(atoms, \mathbb{R})$, can express all additive utility functions.*

Proposition 2 *$\mathcal{L}(pcubes, \mathbb{R})$, the language of positive cubes (conjunctions of atoms), can express all utility functions.*

Proof: Let $u : 2^G \to \mathbb{R}$ be any utility function. Define a goalbase:

$(\top, w_\top)$ with $w_\top = u(\emptyset)$
(p, w_p) with $w_p = u(\{p\}) - w_\top$
$(p \wedge q, w_{p\wedge q})$ with $w_{p\wedge q} = u(\{p, q\}) - w_p - w_q - w_\top$ ...
$(\bigwedge S, w_{\wedge S})$ with $w_{\wedge S} = u(S) - \sum_{S' \subset S} w_{\wedge S'}$ (for $S \subseteq G$)

This goalbase will generate the function u. ✓

Observe that the proof also demonstrates that $\mathcal{L}(pcubes, \mathbb{R})$ has a *unique* way of representing any given function.

$\mathcal{L}(cubes, \mathbb{R})$, for example, is also fully expressive but *not unique:*
$\{(p \wedge q, 5), (p \wedge \neg q, 5), (\neg p \wedge q, 3), (\neg p \wedge \neg q, 3)\} \equiv \{(\top, 3), (p, 2)\}$

Relative Succinctness

Allowing for negation can improve representational succinctness:

Proposition 3 *The language $\mathcal{L}(cubes, \mathbb{R})$ is strictly more succinct than $\mathcal{L}(pcubes, \mathbb{R})$: it will never need more but sometimes exponentially less space to represent the same utility function.*

Proof: Every pcube is a cube, so $\mathcal{L}(cubes, \mathbb{R})$ cannot be worse.

Now consider function u with $u(\emptyset) = 1$ and $u(S) = 0$ for all $S \neq \emptyset$. The following two languages both generate u:

- $\{(\neg p_1 \wedge \cdots \wedge \neg p_n, 1)\} \in \mathcal{L}(cubes, \mathbb{R})$ has *linear* size
- $\{(\bigwedge S, (-1)^{|S|}) \mid S \subseteq G\} \in \mathcal{L}(pcubes, \mathbb{R})$ has *exponential* size

But there can be not better way of expressing u in pcubes, because we have seen that pcube representations are *unique*. ✓

Summary

Many different options for how to measure the *fairness* of allocations:

> Pareto efficiency, utilitarian social welfare, Nash social welfare, egalitarian social welfare or other k-rank dictatorships, leximin, Lorentz curve, inequality indices, envy-freeness and its relaxations, proportionality or maximin share, and more.

The *axiomatic method* can help us to choose between them.

A related consideration concerns the (implicit) assumptions the choice of our social objective makes regarding the *nature of preferences:*

- Is preference intensity meaningful? Ordinal preferences maybe better?
- Does it make sense to compare utilities across agents?

And in practice we need to fix a language for *representing preferences*, and consider its properties (such as *expressivity* and *succinctness*).

What next? We will switch to algorithmic concerns and discuss several concrete methods for finding a fair allocation.