

Computational Social Choice: Autumn 2012

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

Our next major topic is *fair division*: how should we divide one or several goods amongst two or more agents in a fair manner?

This is a problem of social choice, but:

- In this literature, preferences are usually modelled as *utility functions* (rather than as preference orders).
- Fair division problems have an *internal structure*, absent from most voting problems (exception: voting in combinatorial domains).

Today we will introduce a range of fairness criteria for such problems:

- *social welfare orderings* and *collective utility functions*
- *proportionality*, *envy-freeness*, and *degrees of envy*

This material is also covered in my lecture notes. The material on axiomatic foundations is taken from the excellent book by Moulin (1988).

U. Endriss. *Lecture Notes on Fair Division*. ILLC, University of Amsterdam, 2010.

H. Moulin. *Axioms of Cooperative Decision Making*. CUP, 1988.

Notation and Terminology

- Let $\mathcal{N} = \{1, \dots, n\}$ 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.
- Most criteria will not be specific to allocation problems, so we also speak of *agreements* (or *outcomes*, *solutions*, *alternatives*, *states*).
- Each agent $i \in \mathcal{N}$ has a *utility function* u_i (or *valuation function*), mapping agreements to the reals, to model their preferences.
 - Typically, u_i first defined on bundles, so: $u_i(A) = u_i(A(i))$.
 - Discussion: preference intensity, interpersonal comparison
- An agreement A gives rise to a *utility vector* $(u_1(A), \dots, u_n(A))$.
- Sometimes, we are going to define social preference structures directly over utility vectors $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{R}^n$, rather than speaking about the agreements generating them.

Pareto Efficiency

Agreement A is *Pareto dominated* by agreement A' if $u_i(A) \leq u_i(A')$ 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).

Discussion:

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

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.

Common definition of *social welfare*:

$$SW(\mathbf{u}) = \sum_{i \in \mathcal{N}} u_i$$

That is, social welfare is defined as the sum of the individual utilities. Maximising this function amounts to maximising *average utility*.

This is a reasonable definition, but it does not capture everything . . .

► We need a systematic approach to defining social preferences.

Social Welfare Orderings

A *social welfare ordering* (SWO) \preceq is a binary relation over \mathbb{R}^n that is *reflexive*, *transitive*, and *complete*.

Intuitively, if $u, v \in \mathbb{R}^n$, then $u \preceq v$ means that v is socially preferred over u (not necessarily strictly).

We also use the following notation:

- $u \prec v$ iff $u \preceq v$ but not $v \preceq u$ (*strict social preference*)
- $u \sim v$ iff both $u \preceq v$ and $v \preceq u$ (*social indifference*)

Collective Utility Functions

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

Every CUF *induces* a SWO: $u \preceq v \Leftrightarrow SW(u) \leq SW(v)$

Discussion: It is often convenient to think of SWOs in terms of CUFs, but not all SWOs are representable as CUFs (example to follow).

Utilitarian Social Welfare

One approach to social welfare is to try to maximise overall profit. This is known as classical utilitarianism (advocated, amongst others, by Jeremy Bentham, British philosopher, 1748–1832).

The *utilitarian* CUF is defined as follows:

$$SW_{\text{util}}(\mathbf{u}) = \sum_{i \in \mathcal{N}} u_i$$

So this is what we have called “social welfare” a few slides back.

Remark: We define CUFs and SWOs on utility vectors, but the definitions immediately extend to allocations:

$$SW_{\text{util}}(A) = SW_{\text{util}}((u_1(A), \dots, u_n(A))) = \sum_{i \in \mathcal{N}} u_i(A(i))$$

Egalitarian Social Welfare

The *egalitarian* CUF measures social welfare as follows:

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

Maximising this function amounts to improving the situation of the weakest member of society.

The egalitarian variant of welfare economics is inspired by the work of John Rawls (American philosopher, 1921–2002) and has been formally developed, amongst others, by Amartya Sen since the 1970s (Nobel Prize in Economic Sciences in 1998).

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

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

Utilitarianism vs. Egalitarianism

- In the computer science literature the utilitarian viewpoint (that is, social welfare = sum of individual utilities) is often taken for granted. In philosophy, economics, political science not.
- John Rawls' "*veil of ignorance*" (*A Theory of Justice*, 1971):
|| *Without knowing what your position in society (class, race, sex, ...) will be, what kind of society would you choose to live in?*
- Reformulating the *veil of ignorance for multiagent systems*:
|| *If you were to send a software agent into an artificial society to negotiate on your behalf, what would you consider acceptable principles for that society to operate by?*
- Conclusion: worthwhile to investigate egalitarian (and other) social principles for concrete applications in computer science.

Nash Product

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

$$SW_{\text{nash}}(\mathbf{u}) = \prod_{i \in \mathcal{N}} u_i$$

This is a useful measure of social welfare as long as all utility functions can be assumed to be positive. Named after John F. Nash (Nobel Prize in Economic Sciences in 1994; Academy Award in 2001).

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

Ordered Utility Vectors

For any $\mathbf{u} \in \mathbb{R}^n$, the *ordered utility vector* \mathbf{u}^* is defined as the vector we obtain when we rearrange the elements of \mathbf{u} in increasing order.

Example: Let $\mathbf{u} = (5, 20, 0)$ be a utility vector.

- $\mathbf{u}^* = (0, 5, 20)$ means that the weakest agent enjoys utility 0, the strongest utility 20, and the middle one utility 5.
- Recall that $\mathbf{u} = (5, 20, 0)$ means that the first agent enjoys utility 5, the second 20, and the third 0.

Rank Dictators

The *k-rank dictator* CUF for $k \in \mathcal{N}$ is mapping utility vectors to the utility enjoyed by the k -poorest agent:

$$SW_k(\mathbf{u}) = u_k^*$$

Interesting special cases:

- For $k = 1$ we obtain the *egalitarian* CUF.
- For $k = n$ we obtain an *elitist* CUF measuring social welfare in terms of the happiest agent.
- For $k = \lfloor \frac{n+1}{2} \rfloor$ we obtain the *median-rank-dictator* CUF.

The Leximin Ordering

We now introduce a SWO that may be regarded as a refinement of the SWO induced by the egalitarian CUF.

The *leximin ordering* \preceq_{lex} is defined as follows:

$$\mathbf{u} \preceq_{\text{lex}} \mathbf{v} \Leftrightarrow \mathbf{u}^* \text{ lexically precedes } \mathbf{v}^* \text{ (not necessarily strictly)}$$

That means: $\mathbf{u}^* = \mathbf{v}^*$ or there exists a $k \leq n$ such that

- $u_i^* = v_i^*$ for all $i < k$ and
- $u_k^* < v_k^*$

Example: $\mathbf{u} \prec_{\text{lex}} \mathbf{v}$ for $\mathbf{u}^* = (0, 6, 23, 35)$ and $\mathbf{v}^* = (0, 6, 24, 25)$

Lack of Representability

Not every SWO is representable by a CUF:

Proposition 1 *The **leximin ordering** (when defined on \mathbb{R}^n , the full space of utility vectors) is **not** representable by a **CUF**.*

The proof on the next slide closely follows Moulin (1988). We give the proof for $n = 2$ agents (which easily extends to $n > 2$).

H. Moulin. *Axioms of Cooperative Decision Making*. CUP, 1988.

Proof

Assumption: \exists CUF $SW(u_1, u_2)$ representing the leximin ordering \preceq_{lex} .

Define $\epsilon_x := SW(x, 4) - SW(x, 3)$ for all $x \in [1, 2]$.

- Due to $(x, 3) \prec_{\text{lex}} (x, 4)$, we must have $\epsilon_x > 0$ for all $x \in [1, 2]$.

Define $A(p) = \{x \in [1, 2] \mid \epsilon_x \geq \frac{1}{p}\}$ for each $p \in \mathbb{N}$.

Choose $p_0 \in \mathbb{N}$ such that $A(p_0)$ *infinite* and $1, 2 \in A(p_0)$ (exists!).

For any $x, y \in A(p_0)$ with $x < y$ we have:

- $(x, 4) \prec_{\text{lex}} (y, 3) \Rightarrow SW(x, 4) < SW(y, 3)$ (*)
- $\epsilon_x \geq \frac{1}{p_0} \Rightarrow SW(x, 4) - SW(x, 3) \geq \frac{1}{p_0} \stackrel{(*)}{\Rightarrow} SW(y, 3) - SW(x, 3) \geq \frac{1}{p_0}$ (**)

Now consider a finite sequence $x_1 = 1 < x_2 < \dots < x_K = 2$ in $A(p_0)$:

- We have $\sum_{k=2}^K [SW(x_k, 3) - SW(x_{k-1}, 3)] \stackrel{(**)}{\geq} \frac{K-1}{p_0}$,
- but also $\sum_{k=2}^K [SW(x_k, 3) - SW(x_{k-1}, 3)] = SW(2, 3) - SW(1, 3)$.

This is a contradiction (the sum is both unbounded and a fixed value). \checkmark

Axiomatic Method

So far we have simply defined some SWOs and CUFs and informally discussed their attractive and less attractive features.

Next we give a couple of examples for *axioms* — properties that we may or may not wish to impose on a SWO.

Interesting results are then of the following kind:

- A given SWO may or may not satisfy a given axiom.
- A given (class of) SWO(s) may or may not be the only one satisfying a given (combination of) axiom(s).
- A given combination of axioms may be *impossible* to satisfy.

The Pigou-Dalton Principle

A fair SWO will encourage inequality-reducing welfare redistributions.

Axiom 1 (PD) *A SWO \preceq respects the Pigou-Dalton principle if, for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, $\mathbf{u} \preceq \mathbf{v}$ holds whenever there exist $i, j \in \mathcal{N}$ such that:*

- $u_k = v_k$ for all $k \in \mathcal{N} \setminus \{i, j\}$ — only i and j are involved;
- $u_i + u_j = v_i + v_j$ — the change is mean-preserving; and
- $|u_i - u_j| > |v_i - v_j|$ — the change is 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 enjoy very different utilities before the encounter, it may not be meaningful to use their absolute utilities afterwards to assess social welfare, but rather their *relative* gain or loss 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 \preceq is *zero independent* if $u \preceq v$ entails $(u + w) \preceq (v + w)$ for all $u, v, w \in \mathbb{R}^n$.

Example: The SWO induced by the utilitarian CUF is zero independent, while the egalitarian SWO is not.

In fact, a SWO satisfies ZI *iff* it is represented by the utilitarian CUF. See 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 personal utility using different “currencies”. So a desirable property of a SWO may be to be independent of the utility scales used by individual agents.

Assumption: Here, we use positive utilities only: $\mathbf{u} \in (\mathbb{R}^+)^n$.

Notation: Let $\mathbf{u} \cdot \mathbf{v} = (u_1 \cdot v_1, \dots, u_n \cdot v_n)$.

Axiom 3 (SI) A SWO \preceq over positive utilities is *scale independent* if $\mathbf{u} \preceq \mathbf{v}$ entails $(\mathbf{u} \cdot \mathbf{w}) \preceq (\mathbf{v} \cdot \mathbf{w})$ for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in (\mathbb{R}^+)^n$.

Example: Clearly, neither the utilitarian nor the egalitarian SWO are scale independent, but the Nash SWO is.

By a similar result as the one mentioned before, a SWO satisfies SI *iff* it is represented by the Nash CUF.

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 \preceq is *independent of the common utility pace* if $\mathbf{u} \preceq \mathbf{v}$ entails $f(\mathbf{u}) \preceq f(\mathbf{v})$ for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ and for every increasing bijection $f : \mathbb{R} \rightarrow \mathbb{R}$ (applied component-wise to utility vectors).

For a SWO satisfying ICP only interpersonal comparisons ($u_i \leq v_j$ or $u_i \geq v_j$) matter, but not the (cardinal) intensity of $u_i - v_j$.

Example: The utilitarian SWO is *not* independent of the common utility pace, but the egalitarian SWO is. Any k -rank dictator SWO is.

Proportionality

If utility functions are *monotonic* ($B \subseteq B' \Rightarrow u(B) \leq u(B')$), then agents may want the *full* bundle and feel entitled to $1/n$ of its value.

For monotonic utilities, the following definition makes sense:

An allocation A is *proportional* if $u_i(A(i)) \geq \frac{1}{n} \cdot \hat{u}_i$ for every agent $i \in \mathcal{N}$, where \hat{u}_i is the utility given to the full bundle by agent i .

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

Remark: Mostly used in the context of *additive* utilities.

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:

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

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

Degrees of Envy

As we cannot always ensure envy-free allocations, another approach would be to try to *reduce* envy as much as possible.

But what does that actually mean?

A possible approach to systematically defining different ways of measuring the *degree of envy* of an allocation:

- Envy between two agents:
 $\max\{u_i(A(j)) - u_i(A(i)), 0\}$ or
1 if $u_i(A(j)) > u_i(A(i))$ and 0 otherwise
- Degree of envy of a single agent:
max, sum
- Degree of envy of a society:
max, sum [or indeed any SWO/CUF]

Summary

In preparation for our study of *fair division*, we have switched from ordinal preference relations to *utility functions*.

The quality of an allocation can be measured using a variety of *fairness* and *efficiency* criteria (all of them aggregating individual utilities).

We have seen Pareto efficiency, collective utility functions (utilitarian, Nash, egalitarian and other k -rank dictators), the leximin ordering, proportionality, and envy-freeness.

Moulin (1988) provides an excellent introduction to welfare economics and has more material on the axiomatics of social welfare orderings.

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

What next?

Next we will discuss specific problems in fair division. The goal will always be to find an allocation of goods to agents that is optimal in view of one of the criteria discussed today.

We will distinguish:

- allocation of *indivisible* goods (giving rise to combinatorial optimisation problems)
- allocation of a *divisible* good (specifically: cake-cutting problems)