Introduction to Computational Social Choice

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

Lecture 5

We have already seen several fairness and efficiency criteria for collective agreements, such as Pareto optimality or envy-freeness. The field of *welfare economics* has a more general take on this:

- *social welfare orderings* and *collective utility functions*
- introduction to the *axiomatisation* of social welfare orderings

One application of such criteria is in *multiagent resource allocation*. We will give an introduction to MARA mechanisms for the *allocation of indivisible goods*. Specifically:

- brief mentioning of some *complexity results*
- example for a *convergence result* for a *distributed mechanism*
- outlook on other research questions

Notation

- Let $\mathcal{A} = \{1, \ldots, n\}$ be our society of agents throughout.
- We have to decide on an *agreement*. This may be an allocation of goods, possibly coupled with monetary side payments.
- Each agent i has a *utility function* u_i over alternative agreements (which also induces a *preference ordering* $\preceq_i$).
- An agreement x gives rise to a *utility vector* $\langle u_1(x), \ldots, u_n(x) \rangle$
- Often, we can define social preference structures directly over utility vectors $u = \langle u_1, \ldots, u_n \rangle$ (elements of $\mathbb{R}^n$), rather than speaking about the agreements generating them.

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 $W : \mathbb{R}^n \to \mathbb{R}$ mapping utility vectors to the reals.
- Intuitively, if $u \in \mathbb{R}^n$, then $W(u)$ is the utility derived from u by society as a whole.
- Every CUF *represents* an SWO: $u \preceq v \Leftrightarrow W(u) \leq W(v)$

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:

$$sw_u(u) \quad = \quad \sum_{i \in \mathcal{A}gents} u_i$$

Observe that maximising this function amounts to maximising the *average utility* enjoyed by agents in the system.

Egalitarian Social Welfare

The *egalitarian* CUF measures social welfare as follows:

$$sw_e(u) \quad = \quad min\{u_i \mid i \in \mathcal{A}gents\}$$

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.

Nash Product

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

$$sw_N(u) \quad = \quad \prod_{i \in \mathcal{A}gents} u_i$$

This is a useful measure of social welfare as long as all utility functions are 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 reductions ($2 \cdot 6 < 4 \cdot 4$).

Ordered Utility Vectors

We now need some more notation ...

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

Example: Let $u = \langle 5, 20, 0 \rangle$ be a utility vector.

- $\vec{u} = \langle 0, 5, 20 \rangle$ means that the weakest agent enjoys utility 0, the strongest utility 20, and the middle one utility 5.
- Recall that $u = \langle 5, 20, 0 \rangle$ means that the first agent enjoys utility 5, the second 20, and the third 0.

Rank Dictators

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

$$sw_k(u) \quad = \quad \vec{u}_k$$

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 $k = \lfloor n/2 \rfloor$ we obtain the *median-rank-dictator* CUF.

The Leximin-Ordering

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

The *leximin-ordering* $\preceq_\ell$ is defined as follows:

$$u \preceq_\ell v \;\Leftrightarrow\; \vec{u} \text{ lexically precedes } \vec{v} \text{ (not necessarily strictly)}$$

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

- $\vec{u}_i = \vec{v}_i$ for all $i < k$ and
- $\vec{u}_k < \vec{v}_k$

Example: $u \prec_\ell v$ for $\vec{u} = \langle 0, 6, 20, 29\rangle$ and $\vec{v} = \langle 0, 6, 24, 25\rangle$

Axiomatic Approach

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 an 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).

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 an SWO may be to be independent from what individual agents consider "zero" utility.

Axiom 1 (ZI) *An SWO* $\preceq$ *is zero independent iff* $u \preceq v$ *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 CUF is not.

In fact, an 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 an SWO may be to be independent from the utility scales used by individual agents.

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

Notation: Let $u \cdot v = \langle u_1 \cdot v_1, \ldots, u_n \cdot v_n \rangle$.

Axiom 2 (SI) *An SWO $\preceq$ is scale independent iff $u \preceq v$ entails $(u \cdot w) \preceq (v \cdot w)$ for all $u, v, w \in (\mathbb{R}^+)^n$.*

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

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

Independence of the Common Utility Pace

Another desirable property of an 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 3 (ICP) *An SWO* $\preceq$ *is independent of the common utility pace iff* $u \preceq v$ *entails* $f(u) \preceq f(v)$ *for all* $u, v \in \mathbb{R}^n$ *and for every increasing bijection* $f : \mathbb{R} \to \mathbb{R}$.

For an SWO satisfying ICP, only relative comparisons ($u_i \leq u_j$ vs. $u_i \geq u_j$) matter, but the (cardinal) intensities $u_i - u_j$ don't.

Example: The utilitarian CUF does not satisfy ICP, but the egalitarian CUF does. Any k-rank dictator CUF does.

Other Fairness and Efficiency Criteria

Recall that we have already seen some other criteria for assessing fairness and efficiency of a collective agreement:

- Pareto efficiency
- Proportionality
- Envy-freeness

Allocation of Indivisible Goods

Next we will consider the problem of allocating indivisible goods. We can distinguish two approaches:

- In the *centralised approach* (e.g. combinatorial auctions), we need to devise an optimisation algorithm to compute an allocation meeting our fairness and efficiency requirements.
- In the *distributed approach*, allocations emerge in response to agents implementing a sequence of local deals. What can we say about the properties of these emerging allocations?

Setting

For the remainder of today we will work in this framework:

- Set of *agents* $\mathcal{A} = \{1, \ldots, n\}$; finite set of indivisible *goods* $\mathcal{G}$.
- An *allocation* A is a partitioning of $\mathcal{G}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_5, r_7\}$ — agent i owns resources r_5 and r_7
- Each agent $i \in \mathcal{A}$ has got a *valuation function* $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}$.
  Example: $v_i(A) = v_i(A(i)) = 577.8$ — agent i is pretty happy

Complexity Results

Before we look into the "how", here are some complexity results:

- Checking whether an allocation is *Pareto efficient* is coNP-complete.
- Finding an allocation with maximal *utilitarian* social welfare is NP-hard. If all valuations are *modular* then it is polynomial.
- Finding an allocation with maximal *egalitarian* social welfare is also NP-hard, even when all valuations are modular.
- Checking whether an *envy-free* allocation exists is NP-complete; checking whether a *Pareto efficient envy-free* allocation exists is even Σ_2^p-complete.

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.

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 valuation. A *payment function* is a function $p : \mathcal{A} \to \mathbb{R}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.

  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 resources evolve in terms of social welfare?

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 its individual welfare:

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

That is, an agent will only accept a deal *iff* it results in a gain in value (money) that strictly outweighs any loss in money (value).

The Global/Social Perspective

For now, suppose that as system designers we are interested in maximising *utilitarian social welfare*:

$$sw_u(A) \quad = \quad \sum_{i \in \mathcal{A}gents} v_i(A)$$

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 have the following valuation functions:

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

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 *set* $\{chair, table\}$.

Linking the Local and the Global Perspectives

It turns out that individually rational deals are exactly those deals that increase social welfare:

Lemma 1 (Rationality and social welfare) *A deal $\delta = (A, A')$ with side payments is individually rational iff $sw_u(A) < sw_u(A')$.*

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

"$\Leftarrow$": The social surplus can be divided amongst all deal participants by using, say, the following payment function:

$$p(i) \;=\; v_i(A') \;-\; v_i(A) \;-\; \underbrace{\frac{sw_u(A') - sw_u(A)}{|\mathcal{A}|}}_{>0} \qquad \checkmark$$

<u>Discussion:</u> The lemma confirms that individually rational behaviour is "appropriate" in utilitarian societies.

Termination

We can now prove a first result on negotiation processes:

Lemma 2 (Termination) *There can be no infinite sequence of IR deals; that is, negotiation must always terminate.*

Proof: Follows from the first lemma and the observation that the space of distinct allocations is finite. ✓

Convergence

It is now easy to prove the following *convergence* result (originally stated by Sandholm in the context of distributed task allocation):

Theorem 1 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Proof: Termination has been shown in the previous lemma. So let A be the terminal allocation. Assume A is *not* optimal, i.e. there exists an allocation A' with $sw_u(A) < sw_u(A')$. Then, by our first lemma, $\delta = (A, A')$ is individually rational $\Rightarrow$ contradiction. ✓

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

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

More MARA

- Convergence only works if arbitrarily complex deals are allowed. Can *simple preferences* guarantee convergence by *simple deals*?
- What about convergence for *other social optimality criteria*?
- What about other models (e.g. *sharable goods*, *agents on a graph*)?
- Can we give bounds on the number of deals required to reach the optimum ($\rightsquigarrow$ *communication complexity*)?
- How close can we get to the optimum (and how fast) if full convergence cannot be guaranteed? Maybe *simulation* can help.
- What would be suitable logics for specifying MARA mechanisms and, say, verifying convergence results ($\rightsquigarrow$ *social software*)?

Summary

What we have covered today:

- social welfare orderings and collective utility functions for formalising fairness and efficiency criteria
- a first taste of the "axiomatic method" in welfare economics
- introduction to multiagent resource allocation, specifically distributed mechanisms for allocating indivisible goods

Some remarks in relation to earlier lectures:

- MARA with indivisible goods is an example for collective decision making in combinatorial domains (observe that for cake-cutting the number of alternatives is infinite)
- MARA and fair division problems are more specialised collective decision making problems than voting: "*no externalities*" means that agents will be indifferent between a large number of alternatives (all allocations where they receive the same bundle)

Literature

Moulin (1988) provides an excellent introduction to welfare economics, covering the axiomatics of SWOs in detail.

The "MARA Survey" (Chevaleyre et al., 2006) covers most of the material discussed today, and more.

To find out more about convergence in distributed negotiation you may start by consulting the JAIR-2006 paper cited below.

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

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.

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

Last Slide

I hope I convinced you: COMSOC is an exciting interdisciplinary research area. *This is a good time to get into it.*

Besides the papers cited on the slides, particularly the survey papers, a good source of information are the proceedings of the *COMSOC Workshops* (Amsterdam 2006 and Liverpool 2008).

I teach a full-semester *course* on COMSOC at the ILLC, and you can find more slides, references, and exercises here:

`http://www.illc.uva.nl/~ulle/teaching/comsoc/`

Subscribe to the COMSOC *mailing list* for announcements:

`http://lists.duke.edu/sympa/info/comsoc`