Multiagent Systems: Spring 2006

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

Social Choice and Welfare

- Last week we have had a first look into how we may be able to define *what constitutes a "good" allocation* of resources amongst the agents making up an artificial society (the multiagent system).
- We have also seen that what is needed is a mapping from the set of individual preference structures of the agents to some sort of *"social preference structure"* over alternative allocations.
- This kind of problem has been studied in *Welfare Economics* and *Social Choice Theory* for some time.

Plan for Today

Today will be an introduction to these fields, particularly in as far as they concern the definition and analysis of social preferences.

- Ordinal and cardinal preferences of *individual* agents
- Introduction to the *fairness-efficiency* dilemma
- *Social welfare orderings* and *collective utility functions*

This lecture is largely based on Chapters 1 and 2 of the following book:

- H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.

Abstraction 1: Allocations and Agreements

- Our grand aim in this course is to better understand the *multiagent resource allocation* problem.
- So, at the end of the day, we are interested in modelling individual and social preferences over alternative *allocations*.
- Today we'll just speak about *agreements* between agents.
  - This is more *abstract*, because allocations come with a certain structure: some potential solutions will not be feasible etc.
  - In particular, at this stage we shall not worry about the *combinatorial explosion* we face when representing preferences over the allocation of multiple goods.
- We also say: decisions, outcomes, or simply alternatives.

Ordinal Preferences

- The *preference relation* of agent i over alternative agreements:

$$x \preceq_i y \quad \Leftrightarrow \quad \text{agreement } x \text{ is not better than } y \text{ (for agent } i\text{)}$$

- We shall also use the following notation:
  - $x \prec_i y$ iff $x \preceq_i y$ but not $y \preceq_i x$ (*strict preference*)
  - $x \sim_i y$ iff both $x \preceq_i y$ and $y \preceq_i x$ (*indifference*)
- A preference relation $\preceq_i$ is usually required to be
  - *transitive:* if you prefer x over y and y over z, you should also prefer x over z; and
  - *connected:* for any two agreements x and y, you can decide which one you prefer (or whether you value them equally).
- Discussion: useful model, but not without problems (humans cannot always assign rational preferences . . .)

Utility Functions

- Cardinal (as opposed to ordinal) preference structures can be expressed via utility functions . . .
- A *utility function* u_i (for agent i) is a mapping from the space of agreements to the reals.
- Example: $u_i(x) = 10$ means that agent i assigns a value of 10 to agreement x.
- A utility function u_i *representing* the preference relation $\preceq_i$:

$$x \preceq_i y \quad \Leftrightarrow \quad u_i(x) \leq u_i(y)$$

- Discussion: utility functions are very useful, but they suffer from the same problems as ordinal preference relations — even more so (humans typically do not reason with numerical utilities . . .)

The Unanimity Principle

An agreement x is *Pareto-dominated* by another agreement y iff:

- $x \preceq_i y$ for all members i of society; and
- $x \prec_i y$ for at least one member i of society.

An agreement is *Pareto optimal* (or Pareto *efficient*) iff it is not Pareto-dominated by any other feasible agreement (named so after Vilfredo Pareto, Italian economist, 1848–1923).

The *Unanimity Principle* states that society should not select an agreement that is Pareto dominated by another feasible agreement.

The Equality Principle

"All men are created equal . . ."

Equality is probably the most obvious *fairness* postulate.

The *Equality Principle* states that the agreement selected by society should give equal utility to all agents.

The Equality-Efficiency Dilemma

Of course, the Equality Principle may not always be satisfiable, namely if there exists no feasible agreement giving equal utility to everyone.

But even when there *are* equal outcomes, they may not be compatible with the Unanimity Principle. Example:

> Ann and Bob need to divide four items between them: a piano, a precious vase, an oriental carpet, and a lawn-mower. Ann just wants the piano: she will assign utility 10 to any bundle containing the piano, and utility 0 to any other bundle. Bob only cares about how many items he receives: his utility will be 5 times the cardinality of the bundle he receives . . .

Minimising Inequality

So the *pure* Equality Principle seems too strong ...

Instead, we could try to *minimise inequality*. In the case of two agents, a first idea would be to select the agreement x minimising $|u_1(x) - u_2(x)|$ amongst all Pareto optimal agreements.

Example: Suppose there are two feasible agreements x and y:

$$\begin{array}{rclcrcl} u_1(x) & = & 2 & \quad & u_1(y) & = & 8 \\ u_2(x) & = & 4 & \quad & u_2(y) & = & 3 \end{array}$$

Inequality is lower for x, but y seems "better" (if we swap utilities for y, we get an agreement that would be Pareto-superior to x) ...

▶ There are no easy solutions. We need a systematic approach ...

Abstraction 2: Agreements and Utility Vectors

- Let $\mathcal{A} = \{1, \ldots, n\}$ be our agent society throughout.
- An agreement x gives rise to a *utility vector* $\langle u_1(x), \ldots, u_n(x) \rangle$
- We are going to 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.
- Example: The definition of Pareto-dominance is rephrased as follows. Let $u, v \in \mathbb{R}^n$. Then u is Pareto-dominated by v iff:
  - $u_i \leq v_i$ for all $i \in \mathcal{A}$; and
  - $u_i < v_i$ for at least one $i \in \mathcal{A}$.

Social Welfare Orderings

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

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*)

Terminology: In the (economics) literature, connectedness is usually referred to as "completeness". Furthermore, many authors use the letters R, P and I instead of $\preceq$, $\prec$ and $\sim$.

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: $u \preceq v \Leftrightarrow W(u) \leq W(v)$
- Discussion: often convenient to think of SWOs in terms of a 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 across society. 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.

In the MAS literature, the utilitarian CUF is often regarded as the most important notion of social welfare, but there are also others . . .

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.

Ordered Utility Vectors

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.

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 $\vec{v} = \langle 0, 6, 24, 25 \rangle$

Lack of Representability

Not every SWO is representable by a CUF:

Theorem 1 *The leximin-ordering is not representable by a CUF.*

Proof idea: Derive a contradiction by identifying an unbounded sequence of agreements such that (1) there would have to be a minimum increase in collective utility from one agreement to the next; and (2) the difference in collective utility between the final and the first element of the sequence would have to be finite.

Weak Representability

A CUF W is said to *weakly represent* the SWO $\preceq$ iff $W(u) < W(v)$ entails $u \prec v$ for all $u, v \in \mathbb{R}^n$.

Equivalently: A CUF W weakly represents the SWO $\preceq$ iff $u \preceq v$ entails $W(u) \leq W(v)$ for all $u, v \in \mathbb{R}^n$.

- The egalitarian CUF weakly represents the leximin-ordering.

Axioms

We are now going to go through several *axioms* — properties that we may or may not wish to impose on an SWO.

We'll be interested in the following kinds of results:

- 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).

This is a nice example of the so-called *axiomatic method* . . .

Anonymity and Unanimity

The following two axioms will be imposed on any SWO:

Axiom 1 (ANO) *An SWO $\preceq$ is said to respect anonymity iff u being a permutation of v entails $u \sim v$ (indifference) for all $u, v \in \mathbb{R}^n$.*

Axiom 2 (UNA) *An SWO $\preceq$ is said to respect unanimity iff $u \prec v$ holds whenever u is Pareto-dominated by v for all $u, v \in \mathbb{R}^n$.*

Note: We are phrasing axioms as definitions of properties. An axiom is satisfied iff the defined property holds for the SWO in question.

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 3 (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.

Zero Independence and Utilitarianism

The axiom ZI characterises the same SWO as the utilitarian CUF:

Theorem 2 (d'Aspremont & Gevers, 1977; Kaneko, 1984) *An SWO is zero independent iff it is represented by the utilitarian CUF.*

Proof: see Moulin (1988)

C. d'Aspremont and L. Gevers. *Equity and the Informational Basis of Collective Choice*. Review of Economic Studies, 44(2):199–209, 1977.

M. Kaneko. *Reformulation of the Nash Social Welfare Function for a Continuum of Individuals*. Social Choice and Welfare, 1:33–43, 1984.

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 4 (SI) *An SWO $\preceq$ over positive utilities 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.

Nash Product

- The *Nash collective utility function* 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 can be assumed to be positive.

- Named after John F. Nash (Nobel Prize in Economic Sciences in 1994; Academy Award in 2001).

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

- The Nash CUF is scale independent.

Scale Independence and the Nash CUF

The axiom SI characterises the same SWO as the Nash CUF:

Theorem 3 *An SWO over positive utility vectors is scale independent iff it is represented by the Nash CUF.*

Proof: This can be shown to be a corollary to Theorem 2 (which links ZI and utilitarianism). For any SWO $\preceq$ over $(\mathbb{R}^+)^n$ define $\preceq'$ over $\mathbb{R}^n$:

$$u \preceq' v \quad \Leftrightarrow \quad \langle 2^{u_1}, \ldots, 2^{u_n} \rangle \preceq \langle 2^{v_1}, \ldots, 2^{v_n} \rangle$$

Observe that (1) $\preceq$ is scale independent iff $\preceq'$ is zero independent; and (2) $\preceq$ is represented by the Nash CUF iff $\preceq'$ is represented by the utilitarian CUF. The claim then follows from Theorem 2. $\square$

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 5 (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 interpersonal comparisons ($u_i \leq v_i$ or $u_i \geq v_i$) matter, but no the (cardinal) intensity of $u_i - v_i$.

Example: The utilitarian CUF is not independent of the common utility pace, but the egalitarian CUF is.

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:

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

For $k = 1$ we obtain the egalitarian CUF. For $k = n$ we obtain an elitist CUF measuring social welfare in terms of the agent that is best off.

Theorem 4 (Hammond, 1976; d'Aspremont & Gevers, 1977) *An SWO is independent of the common utility pace iff it is weakly represented by the k-rank dictator CUF for some $k \in \mathcal{A}$.*

Proof: see Moulin (1988)

P. Hammond. *Equity, Arrow's Conditions, and Rawls' Difference Principle*. Econometrica, 44(4), 793–804, 1976.

C. d'Aspremont and L. Gevers. *Equity and the Informational Basis of Collective Choice*. Review of Economic Studies, 44(2):199–209, 1977.

Further Axioms

Axiom 6 (SEP) *An SWO $\preceq$ is separable iff social welfare changes are independent of non-concerned agents; that is, iff $u \preceq v$ entails $(u + w) \preceq (v + w)$ for all $u, v, w \in \mathbb{R}^n$ with $w_i = 0$ whenever $u_i \neq v_i$.*

Notation: Let $e = \langle 1, 1, \ldots, 1 \rangle$ be the unit vector in $\mathbb{R}^n$.

Axiom 7 (ICZ) *An SWO $\preceq$ is independent of the common zero of utility iff $u \preceq v$ entails $(u + \lambda e) \preceq (v + \lambda e)$ for all $u, v \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$.*

Axiom 8 (ICS) *An SWO $\preceq$ over positive utilities is independent of the common utility scale iff $u \preceq v$ entails $\lambda u \preceq \lambda v$ for all $u, v \in (\mathbb{R}^+)^n$ and all $\lambda \in \mathbb{R}^+$.*

The Pigou-Dalton Principle

A further desirable property of an SWO would be to encourage inequality-reducing redistributions of welfare.

Axiom 9 (PD) *An SWO is said to respect the Pigou-Dalton Principle iff, for all $u, v \in \mathbb{R}^n$, $u \preceq v$ holds whenever there exist $i, j \in \mathcal{A}$ such that the following conditions are met:*

- $u_k = v_k$ *for all* $k \in \mathcal{A} \setminus \{i, j\}$ *— only* i *and* j *are involved;*
- $u_i + u_j = v_i + v_j$ *— the deal is mean-preserving; and*
- $|u_i - u_j| > |v_i - v_j|$ *— the deal is inequality-reducing.*

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

Pigou-Dalton and the Egalitarian CUF

The egalitarian CUF respects the Pigou-Dalton Principle.

In fact, the egalitarian CUF is the only k-rank dictator CUF not violating the Pigou-Dalton Principle (why?).

▶ Any SWO that satisfies both ICP and PD is weakly represented by the egalitarian CUF (corollary to Theorem 4).

Applications in Multiagent Systems

- Later on in the course, we'll need SWOs and CUFs to specify what we consider a good allocation of resources.
- What interpretation of the term *social welfare* is appropriate depends on the application.
- We are going to have to take back our two abstractions. SWOs and CUFs will be defined directly over alternative *allocations* rather than over alternative *utility vectors*.
- For instance, if utilities are defined over bundles of resources, and the bundle agent i receives in allocation A is $A(i)$, then the utilitarian social welfare of allocation A is defined as follows:

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

Summary

- Individual preferences: *utility functions* (or just ordinal relations)
- *Social welfare orderings* and *collective utility functions*
  - general *definition* of the concepts
  - specific SWO/CUFs: *utilitarian*, *egalitarian*, *leximin*, . . .
  - *representability:* leximin-ordering not representable by CUF
  - *axiomatic approach:* axioms characterising SWO/CUFs
- Social Choice and Welfare is a big area. We have concentrated on formalising the relationship between individual preferences and preferences of society as a whole.

  Social *choice* is then about how to make a collective decision and how such decisions relate to the social preferences. Some of this will come up later in the course; some of it will be discussed in Eric Pacuit's course (Arrow's Theorem, Voting Theory).

References

- H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988. Chapters 1 and 2.
- H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- Classic: A.K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.
- Y. Chevaleyre *et al*. Issues in Multiagent Resource Allocation. *Informatica*, 2006. To appear. Section on Social Welfare.

Summary of Axioms

(ANO)	anonymity
(UNA)	unanimity
(ZI)	zero independence
(SI)	scale independence
(ICP)	independence of the common utility pace
(SEP)	separability
(ICZ)	independence of the common zero of utility
(ICS)	independence of the common utility scale
(PD)	Pigou-Dalton principle

What next?

Today we have discussed how we can define what is good for society. In other words, we have discussed *normative* issues:

> *"Accepting this or that axiom, this is what we would like society to end up with".*

But being autonomous, agents may not *want* to cooperate so as to achieve a state of the world that would be considered optimal according to one of our SWOs.

This is where *Game Theory* comes into play ...