What can AI bring to social choice?

Jérôme Lang
IRIT, Université Paul Sabatier
Toulouse, France

1. Introduction
2. Compact preference representation
3. Vote and aggregation
4. Fair division
5. Conclusion

SAIS, Umeå, May 12, 2006

Social choice

Designing and evaluating methods of collective decision making

1. a finite *set of agents* $\mathcal{A} = \{1, ..., n\}$;
2. a [finite] *set of alternatives* $\mathcal{X}$;
3. a *collective preference profile* = a *preference structure* on $\mathcal{X}$ for each agent
4. $\mathcal{P}^n$ set of all collective profiles.

Social choice function $F : \mathcal{P}^n \to \mathcal{X}$

$F(P_1, \ldots, P_n)$ = socially preferred alternative

Aggregation function $H : \mathcal{P}^n \to \mathcal{P}$

$H(P_1, \ldots, P_n)$ = collective preference relation

Social choice correspondence $C : \mathcal{P}^n \to 2^{\mathcal{X}} \setminus \{\emptyset\}$

$C(P_1, \ldots, P_n)$ = set of socially preferred alternatives

Preference structure: mathematical object modelling the preferences of an agent with respect to a set of alternatives (candidates, allocations etc.) $\mathcal{X}$:

- *cardinal preferences*:
  - *numerical preferences* $u : \mathcal{X} \to \mathbb{R}$
  - *qualitative preferences* $u : \mathcal{X} \to V$ qualitative ordered scale
- *ordinal preferences*: R (or $\succeq$) preference relation (transitive + reflexive) on $\mathcal{X}$

Some subdomains of social choice

- *Vote*: agents (*voters*) express their preferences on a set of alternatives (*candidates*) and must come up to choose a candidate (or a nonempty subset of candidates).
- *Resource allocation* (fair division, auctions...): agents express their preferences over combinations of resources they may receive and an allocation must be found.
- ...

Social choice theory

- designing and evaluating formal methods of collective decision making

  Typical results: *impossibility/possibility theorems*

  There exists / there does not exist a social choice procedure meeting requirements (R1),...,(Rp)

  Example: Arrow's theorem

  Any aggregation function defined on all profiles and satisfying unanimity and independence from irrelevant alternatives is dictatorial

- *computational issues are neglected*

  Knowing that a given procedure *can* be computed is generally enough.

AI and social choice theory: two research areas

From social choice theory to AI

importing concepts and procedures from social choice for solving problems arising in AI applications

- societies of artificial agents (voting, negotiating / bargaining, ...)
- aggregation procedures for web site ranking and information retrieval
- vote procedures for clustering and pattern recognition

From AI to social choice theory

using AI notions and algorithms for solving complex group decision making problems.

1. Introduction
2. **Compact preference representation**
3. Vote and aggregation
4. Fair division
5. Conclusion

Key question: *structure* of the set $\mathcal{X}$ of alternatives?

Example 1 choosing a president:

$\mathcal{X} = \{\texttt{John Kerry, George Bush, Ralph Nader}\}$

$\mapsto$ **reasonable size**

Example 2 choosing a common menu:

$\mathcal{X} = \{\texttt{asparagus risotto, foie gras}\} \times \{\texttt{roasted chicken, vegetable curry}\} \times \{\texttt{white wine, red wine}\}$

Example 3 recruiting committee (3 positions, 6 candidates):

$\mathcal{X} = \{A \mid A \subseteq \{a, b, c, d, e, f\},\ |A| \leq 3\}$;

Example 4 allocation of indivisible goods:

$O = \{o_1, \ldots, o_p\}$ set of indivisible goods;

$A = \{a_1, \ldots, a_n\}$ set of agents;

$\mathcal{X} = A^O$.

Key question: *structure* of the set $\mathcal{X}$ of alternatives?

In Examples 2-5,

$$\mathcal{X} = D_1 \times ... \times D_n$$

combinatorial domain

“Explicit representation” (u or $\succeq$ in extenso) *unreasonable*

$\Rightarrow$ languages for **compact** representation of preference

An example: CP-nets [Boutilier, Brafman, Hoos and Poole, 99]

Idea: exploit *conditional preferential independence* between variables

$\{U, V, W\}$ partition of $\mathcal{V}$.

$D_U = \times_{X_i \in U} D_i$ etc.

U is preferentially independent of V (given W) iff

$$\text{for all } u, u' \in Dom(U),\ v, v' \in Dom(V),\ w \in Dom(W),$$
$$(u, v, w) \succeq (u', v, w) \text{ if and only if } (u, v', w) \succeq (u', v', w)$$

An example: CP-nets [Boutilier, Brafman, Hoos and Poole, 99]

$z_1 \succ z_2 \succ z_3$

Z

$\downarrow$

X $\rightarrow$ **Y** $\rightarrow$ **T**

$x \succ \bar{x}$

$x : y \succ \bar{y}$

$\bar{x} : \bar{y} \succ y$

$y \vee z : t_1 \succ t_2 \succ t_3$

$\neg(y \vee z) : t_2 \succ t_3 \succ t_1$

$x : y \succ \bar{y}$

- if $X = x$
- then $Y = y$ preferred to $Y = \bar{y}$
- everything else being equal (*ceteris paribus*)

An example: CP-nets [Boutilier, Brafman, Hoos and Poole, 99]

$\bar{x} \succ x$
$x : \bar{y} \succ y$
$\bar{x} : y \succ \bar{y}$

$\Rightarrow \bar{x}y \succ \bar{x}\bar{y} \succ x\bar{y} \succ xy$

Simplifying assumption: preferences are strict linear orders (no indifference, no incomparability)

$$\langle \succ_1, \ldots, \succ_n \rangle \mapsto x \in \mathcal{X}$$

deterministic voting rule

Some families of voting rules: *positional scoring rules*

- N voters, p candidates
- fixed list of p integers $s_1 \geq \ldots \geq s_p$
- voter i ranks candidate x in position $j \Rightarrow score_i(x) = s_j$
- choose the candidate maximizing $s(x) = \sum_{i=1}^{n} score_i(x)$

Examples:

- $s_1 = 1,\ s_2 = \ldots = s_p = 0 \Rightarrow$ *plurality* rule;
- $s_1 = p - 1,\ s_2 = p - 2,\ \ldots s_p = 0 \Rightarrow$ *Borda* rule;

Some families of voting rules: *Condorcet-consistent rules*

Condorcet winner

$$x \text{ such that } \forall y \neq x, \#\{i, x \succ_i y\} > \#\{i, y \succ_i x\}$$

- the existence of a Condorcet winner is not guaranteed;
- when a Condorcet winner exists, it is unique

A *Condorcet-consistent rule* elects the Condorcet winner when there is one.

Some families of voting rules: *Condorcet-consistent rules*

Examples:

- *Simpson rule* (or *maximin*):
  let $N(x, y)$ be the number of voters who prefer x to y.
  Simpson score: $S(x) = \min_{y \neq x} N(x, y)$
  Simpson winners = candidates maximizing S.

- *Copeland rule*:
  $x >_{maj} y$: a strict majority of voters prefers x to y.
  $C(x) = \#\{y | x >_{maj} y\} - \#\{y | y >_{maj} x\}$
  Copeland winners = candidates maximizing C.

Computing voting rules: small domains

- most voting rules are linear or quadratic
- a few voting rules are NP-hard; example:

  Dodgson for any $x \in \mathcal{X}$, $D(x)$ = smallest number of elementary changes needed to make x a Condorcet winner.

  elementary change = exchange of adjacent candidates in a voter's ranking

  Deciding whether x is a Dodgson winner requires a logarithmic number of calls to NP oracles: $\Delta_2^{\mathsf{P}}(O(\log n))$-complete [Hemaspaandra, Hemaspaandra & Rothe, 97]

Computing voting rules: combinatorial domains

$\mathcal{V} = \{X_1, \ldots, X_n\}$ set of variables

$\mathcal{X} = D_1 \times ... \times D_n$ set of alternatives

D_i value domain for variable X_i.

Naive formulation: given a profile $(\succ_1, ..., \succ_n)$ and a voting rule F, compute $F(\succ_1, \ldots, \succ_N)$.

Major problem: the explicit specification of each R_i has an exponential size (in the number of variables).

it is not reasonable to ask the voters to specify their preferences in an explicit way $\Rightarrow$ compact preference representation + elicitation procedures needed!

Computing voting rules: combinatorial domains

Voting separately on each variable

$\Rightarrow$ *multiple election paradoxes* (Brams, Kilgour & Zwicker 98)

S: build a new swimming pool; T; build a new tennis court.
Suppose the true preferences are

voters 1 and 2	$S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T} \succ ST$
voters 3 and 4	$\bar{S}T \succ S\bar{T} \succ \bar{S}\bar{T} \succ ST$
voter 5	$ST \succ S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T}$

Problem 1: How can voters 1-4 report their projected preference on $\{S, \bar{S}\}$ and $\{T, \bar{T}\}$?

Does voter 1 prefer S to $\bar{S}$ or $\bar{S}$ to S?

Computing voting rules: combinatorial domains

Voting separately on each variable

⇒ *multiple election paradoxes* (Brams, Kilgour & Zwicker 98)

S: build a new swimming pool; T; build a new tennis court.
Suppose the true preferences are

voters 1 and 2	$S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T} \succ ST$
voters 3 and 4	$\bar{S}T \succ S\bar{T} \succ \bar{S}\bar{T} \succ ST$
voter 5	$ST \succ S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T}$

Problem 1: How can voters 1-4 report their projected preference on $\{S, \bar{S}\}$ and $\{T, \bar{T}\}$?

Does voter 1 prefer S to $\bar{S}$ or $\bar{S}$ to S?

⇒ **decision under uncertainty!**

Computing voting rules: combinatorial domains

Voting separately on each variable

⇒ *multiple election paradoxes* (Brams, Kilgour & Zwicker 98)

S: build a new swimming pool; T; build a new tennis court.
Suppose the true preferences are

voters 1 and 2	$S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T} \succ ST$
voters 3 and 4	$\bar{S}T \succ S\bar{T} \succ \bar{S}\bar{T} \succ ST$
voter 5	$ST \succ S\bar{T} \succ \bar{S}T \succ \bar{S}\bar{T}$

Problem 2: suppose they do so by an “optimistic” projection:

- voters 1, 2 and 5: S; voters 3 and 4: $\bar{S}$ ⇒ decision = S;
- voters 3,4 and 5: T; voters 1 and 2: $\bar{T}$ ⇒ decision = T.

Alternative ST is chosen although it is the worst alternative for all but one voter.

Computing voting rules: combinatorial domains

Idea: exploit *conditional preferential independence* between variables

Assumption: voters' preferences share the same preferential independence structure

$\Rightarrow$ they are expressible by CP-nets with a common acyclic graph $\mathcal{G}$.

Example: r = plurality rule

3 voters	2 voters	2 voters
$\bar{x}y \succ \bar{x}\bar{y} \succ x\bar{y} \succ xy$	$xy \succ x\bar{y} \succ \bar{x}\bar{y} \succ \bar{x}y$	$x\bar{y} \succ xy \succ \bar{x}y \succ \bar{x}\bar{y}$
$\Downarrow$	$\Downarrow$	$\Downarrow$
$\bar{x} \succ x$	$x \succ \bar{x}$	$x \succ \bar{x}$
$x : \bar{y} \succ y$	$x : y \succ \bar{y}$	$x : \bar{y} \succ y$
$\bar{x} : y \succ \bar{y}$	$\bar{x} : \bar{y} \succ y$	$\bar{x} : y \succ \bar{y}$

For all voters, X is preferentially independent of Y: $\mathcal{G} = \{(X, Y)\}$ acyclic

1. 4 voters out of 7 unconditionally prefer x over $\bar{x} \Rightarrow x^* = x$;
2. given $X = x$, 5 voters out of 7 prefer $\bar{y}$ to $y \Rightarrow y^* = \bar{y}$;

Sequential plurality winner $= (x, \bar{y})$

Direct plurality winner $= (\bar{x}, y)$

$\Rightarrow$ *plurality is not decomposable*

Computing voting rules: combinatorial domains

A voting rule r on $\mathcal{X} = D_1 \times \ldots \times D_p$ is **decomposable**
iff there exist n voting rules $r_1, \ldots, r_p$ on $D_1, \ldots, D_p$ such that:

for any linear order $\mathcal{O} = X_1 > ... > X_p$ on $\mathcal{V}$
and for any preference profile $R = (R_1, ..., R_N)$ following $\mathcal{O}$,
we have $Seq(r_1, \ldots, r_p)(R) = r(R)$.

- plurality is not decomposable;
- more generally: no positional scoring rule is decomposable

Question: *are there any decomposable rules other than dictatorial rules?*

Sequential Condorcet winners

- if preference on X is unconditional:
  $\vec{x} \in D_X$ is a X-Condorcet winner if and only if

$$(\forall \vec{y} \in D_{\bar{X}}) \forall \vec{x}' \in D_X \ \#\{i, \vec{x}\vec{y} \succ_i \vec{x}'\vec{y}\} > \frac{N}{2}$$

- if preference on Y given $X = \vec{x}$ is unconditional:
  $\vec{y} \in D_Y$ is a Y-Condorcet winner given $X = \vec{x}$ if and only if

$$(\forall \vec{z} \in D_{\bar{X \cup Y}}) \forall \vec{y}' \in D_X \ \#\{i, \vec{x}\vec{y}\vec{z} \succ_i \vec{x}\vec{y}'\vec{z}\} > \frac{N}{2}$$

Sequential CW = sequential combination of “local” CWs

Question: $\vec{s}$ sequential CW $\Rightarrow$ $\vec{s}$ (direct) CW?

Sequential Condorcet winners

- any direct Condorcet winner is a sequential Condorcet winner
- the converse does not hold

2 voters	1 voter	2 voters
$x\bar{y} \succ \bar{x}\bar{y} \succ xy \succ \bar{x}y$	$xy \succ x\bar{y} \succ \bar{x}y \succ \bar{x}\bar{y}$	$\bar{x}y \succ \bar{x}\bar{y} \succ xy \succ x\bar{y}$

X and Y are preferentially independent $\Rightarrow$ take any order

- 3 voters unconditionally prefer x to $\bar{x}$ $\Rightarrow$ x local Condorcet winner
- 3 voters unconditionally prefer y to $\bar{y}$ $\Rightarrow$ y local Condorcet winner

$\Rightarrow$ xy sequential Condorcet winner
but xy is not a direct Condorcet winner (4 voters: $\bar{x}\bar{y} \succ xy$)

paradox known as the *multiple referendums paradox*

Equivalence obtained if preferences are lexicographic. *Other restrictions?*

Preference aggregation: combinatorial domains

Aggregation of compactly represented preferences?

$$\begin{array}{ccc} & H? & \\ \langle \Phi_1, \ldots, \Phi_n \rangle & --\rightarrow & \Phi^* \\ \downarrow & & \downarrow \\ \langle R_1, \ldots, R_n \rangle & \longrightarrow & R^* \\ & h & \end{array}$$

- $\Phi_1, \ldots, \Phi_n$ compact descriptions of preference relations $R_1, \ldots, R_n$
- Φ^* compact description of $h(R_1, \ldots, R_n)$:
  - Φ^* belongs to the same language as $\Phi_1, \ldots, \Phi_n$;
  - Φ^* should be computed directly;
  - Φ^* should be as small as possible

A first step: multiagent CP-nets in [Rossi, Venable & Walsh 04]

Preference aggregation: combinatorial domains

Merging logical belief/preference bases

$$\begin{array}{ccc} & ? & \\ \langle B_1, \ldots B_n \rangle & - - \rightarrow & B^* \\ \downarrow & & \downarrow \\ d(., B) = \star_i d(., B_i) & \longrightarrow & \begin{array}{c}\Delta(B_1, \ldots, B_n) = \\ \{w \text{ minimizing } d(w, B)\}\end{array} \end{array}$$

- $B_1, \ldots, B_n$ [sets/multisets of] propositional formulae
- B^* compact description of a *dichotomous* preference
- interpersonal comparison of preferences needed
- direct computation of B^*: works in some cases

Preference merging and social choice: [Konieczny & Pino-Perez 05; Benferhat, Dubois, Kaci & Prade 02, Chopra, Ghose &Meyer 05]

Manipulation and strategyproofness

Manipulation: a coalition of voters expressing an insincere preference profile so as to give more chance to a preferred candidate to be elected.

Example: $r =$ plurality rule

3 voters	2 voters	2 voters
$\bar{x}y$	xy	$x\bar{y}$
$\bar{x}\bar{y}$	$x\bar{y}$	xy
$x\bar{y}$	$\bar{x}\bar{y}$	$\bar{x}y$
xy	$\bar{x}y$	$\bar{x}\bar{y}$

Outcome: $\bar{x}y$

3 voters	**2 voters**	2 voters
$\bar{x}y$	$x\bar{y}$	$x\bar{y}$
$\bar{x}\bar{y}$	xy	xy
$x\bar{y}$	$\bar{x}\bar{y}$	$\bar{x}y$
xy	$\bar{x}y$	$\bar{x}\bar{y}$

Outcome: $x\bar{y}$

Manipulation and strategyproofness

Gibbard (73) and Satterthwaite (75) 's theorem: if the number of candidates is at least 3, then any nondictatorial voting procedure is manipulable for some profiles.

Barriers to manipulation:

- making manipulation *less efficient*: make as little as possible of the others' votes known to the would-be manipulating coalition
- make manipulation *hard to compute* [Bartholdi, Tovey & Trick, 89]; [Bartholdi & Orlin, 91]; [Conitzer & Sandholm, 02, 03]; [Conitzer, Lang & Sandholm, 03]

Making manipulation computationally hard

$\{(1, \alpha_1), \ldots, (n, \alpha_n)\}$ set of *weighted* voters ($\alpha_i \in \mathbb{N}^*$ for all i)

CONSTRUCTIVE MANIPULATION EXISTENCE: given a voting rule r, a set of p candidates $\mathcal{X}$, a candidate $x \in \mathcal{X}$, and the preferences rankings of voters $1, \ldots, k < n$, is there a way for voters $K + 1, \ldots, n$ to cast their votes such that x is elected?

- plurality: in P;
- all other scoring rules (including Borda and veto): in P for $p = 2$, NP-complete for $p \geq 3$;
- Copeland and Simpson: in P for $p \leq 3$, NP-complete for $p \geq 4$;

[Conitzer & Sandholm, 02]; [Conitzer, Lang & Sandholm, 03]

Incomplete knowledge and communication complexity

Given some *incomplete* description of the voters' preferences,

- is the outcome of the voting rule determined?
- if not, whose information about which candidates is needed?

4 voters: $c \succ d \succ a \succ b$

2 voters: $a \succ b \succ d \succ c$

2 voters: $b \succ a \succ c \succ d$

1 voter: $? \succ ? \succ ? \succ ?$

plurality winner already known (c)

Borda

partial scores (for 8 voters): a: 14 ; b: 10 ; c: 14; d: 10

$\Rightarrow$ only need to know the last voters's preference between a and c

general study in [Conitzer & Sandholm, 02]

Incomplete knowledge and communication complexity

Communication complexity [Yao 79]: measure the minimum amount of information to be communicated so that the outcome of the voting procedure is determined.

⇒ design protocols for gathering the information as economically as possible

Incomplete knowledge and communication complexity

Example: plurality with runoff, n voters, p candidates.

Optimal protocol:

step 1 voters send the name of their most preferred candidate to the central authority C

$\hookrightarrow$ $\mathbf{n} \log \mathbf{p}$ bits

step 2 C sends the names of the two finalists to the voters

$\hookrightarrow$ $\mathbf{2n} \log \mathbf{p}$ bits

step 3 voters send the name of their preferred finalist to C

$\hookrightarrow$ $\mathbf{n}$ bits

total $\mathbf{n}(\mathbf{3} \log \mathbf{p} + \mathbf{1})$ bits (in the worst case)

[Conitzer & Sandholm, 05]

1. Introduction
2. Compact preference representation
3. Vote and aggregation
4. **Fair division**
5. Conclusion

Resource allocation / fair division

$\mathcal{A} = \{1, \ldots, n\}$ agents

$\mathcal{R} = \{r_1, \ldots, r_p\}$ *indivisible* resources (objects)

$\pi : \mathcal{A} \rightarrow 2^{\mathcal{R}}$ *allocation*

Possible requirements for allocations:

- $\pi(i) \cap \pi(j) = \emptyset$ for $i \neq j$: *preemptive allocations*;
- $\cup_i \pi(i) = \mathcal{R}$: *complete allocations*;
- $\pi(i) = \pi(j)$ for all i, j: *shared allocations*

Finding an allocation
= group decision making with a combinatorial set of alternatives

Resource allocation $\neq$ fair division

Combinatorial auctions

$V_i : 2^{\mathcal{R}} \rightarrow I\!N$ for each agent i

$V_i(X)$ maximal value (price) that i is ready to pay for the combination of resources X

V_i additive for all $i \Rightarrow$ elicitation and optimal allocation are easy

V_i generally *not additive*

{left shoe}	5 $	{beer}	4 $
{right shoe}	5 $	{lemonade}	3 $
{left shoe, right shoe}	40 $	{beer, lemonade}	5 $
complementarity (superadditivity)		supplementarity (subadditivity)	

Resource allocation $\neq$ fair division

Combinatorial auctions: given $V_i : 2^{\mathcal{R}} \rightarrow I\!N$ for each agent i,

find the allocation maximizing the seller's revenue:

$$\pi^* \ maximizing \ \sum_{i=1}^{n} V(\pi(i))$$

purely utilitarianistic criterion (no equity/fairness involved)

Computational issues:

- representation / elicitation of the value functions⇒ bidding languages [Sandholm 99; Nisan 00; Boutilier & Hoos 01]
- computation of the optimal allocation (NP-hard): a huge literature

Fair division: three families of criteria

Numerical criteria

Need *numerical preferences* (sums of utilities are meaningful)

- utilitarianism + monetary compensation

agents	1	2
$\{a, b, c\}$	10	10
$\{a, b\}$	8	9
$\{a, c\}$	8	6
$\{b, c\}$	5	5
$\{a\}$	5	4
$\{b\}$	5	3
$\{c\}$	2	4
$\emptyset$	0	0

Fair division: three families of criteria

Numerical criteria

Need *numerical preferences* (sums of utilities are meaningful)

- utilitarianism + monetary compensation

agents	1	2
$\{a, b, c\}$	10	10
$\{a, b\}$	$\boxed{8}$	9
$\{a, c\}$	8	6
$\{b, c\}$	5	5
$\{a\}$	5	4
$\{b\}$	5	3
$\{c\}$	2	$\boxed{4}$
$\emptyset$	0	0

optimal allocation: $\pi = \langle \{a, b\}, \{c\} \rangle$

+ monetary compensation from 1 to 2: $\frac{8-4}{2} = 2$

Fair division: three families of criteria

Qualitative criteria

Need *(at least) qualitative preferences* $u_i : 2^{\mathcal{R}} \to L$ totally ordered scale common to all agents $\Rightarrow$ interpersonal comparison of preference allowed.

Fair division: three families of criteria

Qualitative criteria

Need *(at least) qualitative preferences* $u_i : 2^{\mathcal{R}} \to L$

- **equity** (or egalitarianism): the *leximin* ordering

agents	1	2
$\{a, b, c\}$	10	10
$\{a, b\}$	8	9
$\{a, c\}$	8	6
$\{b, c\}$	5	5
$\{a\}$	5	4
$\{b\}$	5	3
$\{c\}$	2	4
$\emptyset$	0	0

Fair division: three families of criteria

Qualitative criteria

Need *(at least) qualitative preferences* $u_i : 2^{\mathcal{R}} \rightarrow L$ totally ordered scale

- **equity** (or egalitarianism): the *leximin* ordering

agents	1	2
$\{a, b, c\}$	10	10
$\{a, b\}$	8	9
$\{a, c\}$	8	$\boxed{6}$
$\{b, c\}$	5	5
$\{a\}$	5	4
$\{b\}$	$\boxed{5}$	3
$\{c\}$	2	4
$\emptyset$	0	0

optimal allocation:

$\pi = \langle \{b\}, \{a, c\} \rangle$

Fair division: three families of criteria

Ordinal criteria need *(at least) ordinal preferences*

$\geq_i$: $2^{\mathcal{R}} \rightarrow L$ complete preference relation on $2^{\mathcal{R}}$

- **Pareto efficiency**: π is *efficient* iff there is no π' such that $\pi'(i) \geq_i \pi(i)$ for all i and $\pi'(i) >_i \pi(i)$ for at least one i.
- **envy-freeness**: π is *envy-free* iff for all $i, j \neq i$, $\pi(i) \geq_i \pi(j)$

- **Pareto efficiency**: π is *efficient* iff there is no π' such that $\pi'(i) \geq_i \pi(i)$ for all i and $\pi'(i) >_i \pi(i)$ for at least one i.
- **envy-freeness**: π is *envy-free* iff for all $i, j \neq i$, $\pi(i) \geq_i \pi(j)$

agents	1	2
$\{a,b,c\}$	10	10
$\{a,b\}$	8	9
$\{a,c\}$	8	$\boxed{6}$
$\{b,c\}$	5	5
$\{a\}$	5	4
$\{b\}$	$\boxed{5}$	3
$\{c\}$	2	4
$\emptyset$	0	0

$\pi = \langle \{b\}, \{a,c\} \rangle$ Pareto-efficient

but not envy-free: 1 envies 2

- **Pareto efficiency**: π is *efficient* iff there is no π' such that $\pi'(i) \geq_i \pi(i)$ for all i and $\pi'(i) >_i \pi(i)$ for at least one i.
- **envy-freeness**: π is *envy-free* iff for all $i, j \neq i$, $\pi(i) \geq_i \pi(j)$

agents	1	2
$\{a,b,c\}$	10	10
$\{a,b\}$	8	9
$\{a,c\}$	8	6
$\{b,c\}$	5	$\boxed{5}$
$\{a\}$	$\boxed{5}$	4
$\{b\}$	5	3
$\{c\}$	2	4
$\emptyset$	0	0

$\pi' = \langle\{a\},\{b,c\}\rangle$ envy-free but not Pareto-efficient

For this example there is no allocation being both efficient and envy-free

Fair division: three families of criteria

preferences	numerical	qualitative	ordinal
	$u_i : 2^{\mathcal{R}} \rightarrow I\!N$	$u_i : 2^{\mathcal{R}} \rightarrow L$	$\geq_i$ on 2^R
		L ordered scale	
monetary compensations	+	-	-
interpersonal comparisons	+	+	-
intrapersonal comparisons	+	+	+
	utilitarianism	equity	Pareto efficiency envy-freeness

Resource allocation / fair division

- social choice theory: *axiomatic study of criteria*
- AI & OR: computational and representation issues, mainly for combinatorial auctions

⇒ Representation and computational issues for fair division?

– approximate envy-freeness: [Lipton-Markakis-Mossel-Saberi 04]

– logical representation + complexity results for

  – ordinal fair division: [Bouveret Lang 05]

  – cardinal fair division [Bouveret Fargier Lang Lemaître 05]

– complexity issues in *distributed* allocation: [Dunne, Wooldridge Laurence 05; Chevaleyre, Endriss, Estivié Maudet 04]

Fair division under dichotomous preferences [Bouveret Lang, 05]

dichotomous preference relations R is dichotomous if and only if there is a set of "good" bundles *Good* such that for each subsets A, B of $\mathcal{R}$, $A \succeq_R B$ if and only if $A \in Good$ or $B \notin Good$.

Example:

$$X = \{a, b, c\} \Rightarrow 2^X = \{\varnothing, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}\}$$

$$\frac{Good \longrightarrow \{\{a, b\}, \{b, c\}\}}{\overline{Good} \longrightarrow \{\varnothing, \{a\}, \{b\}, \{c\}, \{a, c\}, \{a, b, c\}\}}$$

Fair division under dichotomous preferences

A dichotomous preference is fully defined by its set of good bundles ⇒ *propositional logic representation*

Example:

	Paul (agent 1)	Mary (agent 2)
$Good_i$	$\{\{a,b\},\{b,c\},\{a,b,c\}\}$	$\{\{b\}\{b,c\}\}$
φ_i	$b \wedge (a \vee c)$	$b \wedge \neg a$

$\succeq_{R_i}$ monotonous $\Leftrightarrow$ $Good_i$ upward closed $\Leftrightarrow$ φ_i positive formula

Fair division under dichotomous preferences

Simple propositional representation of the problem

$$\mathcal{P} = \langle \varphi_1, \ldots, \varphi_N \rangle$$

agent i, good $x \mapsto$ propositional variable x_i (x allocated to i)

rewrite φ_i, replacing each x by $x_i \Rightarrow \varphi_i^*$.

Example (continued):

	Paul (agent 1)	Mary (agent 2)
$Good_i$	$\{\{a,b\},\{b,c\},\{a,b,c\}\}$	$\{\{b\}\{b,c\}\}$
φ_i	$b \wedge (a \vee c)$	$b \wedge \neg a$
φ_i^*	$b_1 \wedge (a_1 \vee c_1)$	$b_2 \wedge \neg a_2$

Fair division under dichotomous preferences

allocation $\approx$ truth assignment of the x_i, satisfying:

$$\Gamma_{\mathcal{P}} = \bigwedge_{x \in X} \bigwedge_{i \neq j} \neg(x_i \wedge x_j)$$

Example (continued):

$$\Gamma_{\mathcal{P}} = \neg(a_1 \wedge a_2) \wedge \neg(b_1 \wedge b_2) \wedge \neg(c_1 \wedge c_2)$$

$\pi : [1 \mapsto \{a, c\}, 2 \mapsto \{b\}] \Rightarrow F(\pi) = (a_1, \neg a_2, \neg b_1, b_2, c_1, \neg c_2)$

Fair division under dichotomous preferences

Simple characterization of envy-freeness :

$$\Lambda_{\mathcal{P}} = \bigwedge_{i=1,\ldots,N} \left[\varphi_i^* \vee \left(\bigwedge_{j \neq i} \neg \varphi_{j|i}^* \right) \right]$$

where $\varphi_{j|i}^* = \varphi_i^*(x_i \leftarrow x_j)$

Proposition: π is envy-free if and only if $F(\pi) \models \Lambda_{\mathcal{P}}$.

Example (continued):

$\Lambda_{\mathcal{P}} =$

$$\begin{aligned} & (\overbrace{(b_1 \wedge (a_1 \vee c_1))}^{\text{1 is satisfied with her share}} \vee \overbrace{\neg(b_2 \wedge (a_2 \wedge c_2))}^{\text{1 wouldn't be satisfied with 2's share}}) \\ \wedge\ & (\underbrace{(b_2 \wedge \neg a_2)}_{\text{2 is satisfied with her share}} \vee \underbrace{\neg(b_1 \wedge \neg a_1)}_{\text{2 wouldn't be satisfied with 1's share}}) \end{aligned}$$

Fair division under dichotomous preferences

Pareto-efficiency requires that allocations satisfy a *maximal* set of agents.

Proposition: π is efficient if and only if $\{\varphi_i^* | F(\pi) \vDash \varphi_i^*\}$ is a maximal $\Gamma_{\mathcal{P}}$-consistent subset of $\{\varphi_1^*, \ldots, \varphi_N^*\}$.

Example (continued):

	agent 1	agent 2
$Good_i$	$\{\{a,b\},\{b,c\},\{a,b,c\}\}$	$\{\{b\}\{b,c\}\}$
φ_i	$(b \wedge (a \vee c))$	$b \wedge \neg a$
φ_i^*	$(b_1 \wedge (a_1 \vee c_1))$	$b_2 \wedge \neg a_2$

$$\Gamma_{\mathcal{P}} = \neg(a_1 \wedge a_2) \wedge \neg(b_1 \wedge b_2) \wedge \neg(c_1 \wedge c_2)$$

The 2 maximal $\Gamma_{\mathcal{P}}$-consistent subsets of $\{\varphi_1^*, \varphi_2^*\}$ are $\{\varphi_1^*\}$ and $\{\varphi_2^*\}$

Fair division under dichotomous preferences

Putting things together:

There exists an efficient and envy-free allocation

if and only if

$\exists S$ maximal $\Gamma_{\mathcal{P}}$-consistent subset of $\{\varphi_1^*, \ldots, \varphi_N^*\}$
such that $\bigwedge S \wedge \Gamma_{\mathcal{P}} \wedge \Lambda_{\mathcal{P}}$ is consistent.

$\Rightarrow$ SKEPTICAL INFERENCE IN DEFAULT LOGIC! (Reiter 1980)

Fair division under dichotomous preferences

Definition: Δ a set of formulae, β and ψ formulae.
ψ is a *skeptical consequence* of $\langle \beta, \Delta \rangle$ (denoted $\langle \beta, \Delta \rangle \mid\!\sim^{\forall} \psi$)
iff $\forall S \in MaxCons(\Delta, \beta)$, $\bigwedge S \wedge \beta \models \psi$.

Proposition:

there exists an EEF allocation *iff*

$$\langle \Gamma_{\mathcal{P}}, \{\varphi_1^*, \ldots, \varphi_N^*\} \rangle \not\mid\!\sim^{\forall} \neg \Lambda_{\mathcal{P}}$$

Consequences:

- upper complexity bound;
- using default logic algorithms for finding EEF allocations

Fair division: bipartite fair matching

Two types of agents: $A = \{a_1, \ldots, a_n\}$; $B = \{b_1, \ldots, b_n\}$.

Find a fair matching given preferences of A-agents over B and preferences of B-agents over A

Example: $A = \{a(lice), b(etty), c(harles)\}$; $B = \{Barcelona, London, Paris\}$

a	:	London > Paris > Barcelona	Barcelona	:	a > c > b
b	:	Barcelona > London > Paris	London	:	b > c > a
c	:	London > Barcelona > Paris	Paris	:	c > b > a

Fair division: bipartite fair matching

Example: $A = \{a(lice), b(etty), c(harles)\}$; $B = \{Barcelona, London, Paris\}$

a	:	London > Paris > Barcelona	Barcelona	:	a > c > b
b	:	Barcelona > London > Paris	London	:	b > c > a
c	:	London > Barcelona > Paris	Paris	:	c > b > a

Stable allocation: if candidate x is matched with university u then any university u' such that $u' >_x u$ is matched with a candidate x' such that $x' >_{u'} x$, and similarly for universities.

π_1: a $\mapsto$ Paris, b $\mapsto$ Barcelona, c $\mapsto$ London
π_2: a $\mapsto$ Barcelona, b $\mapsto$ London, c $\mapsto$ Paris

π_1, π_2 stable allocations

Fair division: bipartite fair matching

Example: $A = \{a(lice), b(etty), c(harles)\}$; $B = \{Barcelona, London, Paris\}$

a	:	London > Paris > Barcelona	Barcelona	:	a > c > b
b	:	Barcelona > London > Paris	London	:	b > c > a
c	:	London > Barcelona > Paris	Paris	:	c > b > a

Stable allocation: if candidate x is matched with university u then any university u' such that $u' >_x u$ is matched with a candidate x' such that $x' >_{u'} x$, and similarly for universities.

π_1: a $\mapsto$ Paris, b $\mapsto$ Barcelona, c $\mapsto$ London
π_2: a $\mapsto$ Barcelona, b $\mapsto$ London, c $\mapsto$ Paris

π_1, π_2 stable allocations
π_1 Pareto-efficient for candidates but not for universities
π_2 Pareto-efficient for universities but not for candidates

What I did not tell about

- social software;
- sequential group decision making;
- fairness and uncertainty;
- mechanism design;
- negotiation;
- communication languages;
- ...