A general elicitation-free protocol for allocating indivisible goods

Sylvain Bouveret — Onera Toulouse

Jérôme Lang — Université Paris Dauphine

Twenty second International Joint Conference on Artificial Intelligence

Barcelona, July 16–22, 2011

A fair division problem...

You have...

- p candies of different kinds and flavors;
- n kids: kid A, kid B, kid C...

A fair division problem. . .

You have. . .

- p candies of different kinds and flavors;
- n kids: kid A, kid B, kid C. . .

How would you allocate the candies to the kids so as to be as ***fair*** *as possible ?*

A fair division problem. . .

You have. . .

- p candies of different kinds and flavors;
- n kids: kid A, kid B, kid C. . .

How would you allocate the candies to the kids so as to be as ***fair*** *as possible ?*

- **Answer 1** (the social choice theoretician's): Ask the kids to give their preferences and use a (centralized) collective decision making procedure.
- **Answer 2** (the MAS specialist's): Ask the kids to negotiate.

A fair division problem. . .

You have. . .

- p candies of different kinds and flavors;
- n kids: kid A, kid B, kid C. . .

How would you allocate the candies to the kids so as to be as ***fair*** *as possible ?*

- **Answer 1** (the social choice theoretician's): Ask the kids to give their preferences and use a (centralized) collective decision making procedure.
- **Answer 2** (the MAS specialist's): Ask the kids to negotiate.

A much simpler procedure. . .

A much simpler procedure...

Ask the kids to pick in turn their most preferred candy among the remaining ones, according to some **predefined sequence**.

Example

3 kids A, B, C, 6 candies, sequence $ABCCBA \rightarrow A$ chooses first (and takes her preferred candy), then B, then C, then C again...

A much simpler procedure...

Ask the kids to pick in turn their most preferred candy among the remaining ones, according to some **predefined sequence**.

Example

3 kids A, B, C, 6 candies, sequence $ABCCBA \rightarrow A$ chooses first (and takes her preferred candy), then B, then C, then C again...

We "feel" that $ABCCBA$ is **fairer** than $AABBCC$...

$\rightarrow$ *What is the **fairest** sequence ?*

Agents, objects, preferences...

- A set $\mathcal{O}$ of p objects $\{o_1, \dots, o_p\}$
- A set $\mathcal{N}$ of n agents $\{1, \dots, n\}$
- Each agent i has a (private) **ranking** $\succ_i$ over $\mathcal{O}$ (ex: $o_6 \succ o_1 \succ o_4 \succ o_5 \succ o_2 \succ o_3$)
- A central authority (CA) which must find a **policy** (a sequence of agents) $\pi : \{1, \dots, p\} \to \{1, \dots, n\}$

Example

Example

5 objects, 3 agents, $\pi = 12332\ldots$

- $1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

Example

Example

5 objects, 3 agents, $\pi = 12332\ldots$

- $1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0
$s(1)_k^\pi$	$\emptyset$
$s(2)_k^\pi$	$\emptyset$
$s(3)_k^\pi$	$\emptyset$
$\mathcal{O}_k^\pi$	$\emptyset$

Example

Example

5 objects, 3 agents, $\pi = \mathbf{1}2332\ldots$

- $1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0	1
$s(1)^\pi_k$	$\emptyset$	o_1
$s(2)^\pi_k$	$\emptyset$	$\emptyset$
$s(3)^\pi_k$	$\emptyset$	$\emptyset$
$\mathcal{O}^\pi_k$	$\emptyset$	o_1

Example

Example

5 objects, 3 agents, $\pi = 1\mathbf{2}332\ldots$

- $1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0	1	2
$s(1)_k^\pi$	$\emptyset$	o_1	o_1
$s(2)_k^\pi$	$\emptyset$	$\emptyset$	o_4
$s(3)_k^\pi$	$\emptyset$	$\emptyset$	$\emptyset$
$\mathcal{O}_k^\pi$	$\emptyset$	o_1	$o_1 o_4$

Example

Example

5 objects, 3 agents, $\pi = 12\mathbf{3}32\ldots$

- $1 : o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0	1	2	3
$s(1)^\pi_k$	$\emptyset$	o_1	o_1	o_1
$s(2)^\pi_k$	$\emptyset$	$\emptyset$	o_4	o_4
$s(3)^\pi_k$	$\emptyset$	$\emptyset$	$\emptyset$	o_3
$\mathcal{O}^\pi_k$	$\emptyset$	o_1	$o_1 o_4$	$o_1 o_4 o_3$

Example

Example

5 objects, 3 agents, $\pi = 123\mathbf{3}2\ldots$

- 1 : $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- 2 : $o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- 3 : $o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0	1	2	3	4
$s(1)^\pi_k$	$\emptyset$	o_1	o_1	o_1	o_1
$s(2)^\pi_k$	$\emptyset$	$\emptyset$	o_4	o_4	o_4
$s(3)^\pi_k$	$\emptyset$	$\emptyset$	$\emptyset$	o_3	$o_3 o_5$
$\mathcal{O}^\pi_k$	$\emptyset$	o_1	$o_1 o_4$	$o_1 o_4 o_3$	$o_1 o_4 o_3 o_5$

Example

Example

5 objects, 3 agents, $\pi = 1233\mathbf{2}\ldots$

- 1 : $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- 2 : $o_4 \succ o_2 \succ o_5 \succ o_1 \succ o_3$
- 3 : $o_1 \succ o_3 \succ o_5 \succ o_4 \succ o_2$

k	0	1	2	3	4	5
$s(1)^\pi_k$	$\emptyset$	o_1	o_1	o_1	o_1	o_1
$s(2)^\pi_k$	$\emptyset$	$\emptyset$	o_4	o_4	o_4	$o_4\,o_2$
$s(3)^\pi_k$	$\emptyset$	$\emptyset$	$\emptyset$	o_3	$o_3\,o_5$	$o_3\,o_5$
$\mathcal{O}^\pi_k$	$\emptyset$	o_1	$o_1\,o_4$	$o_1\,o_4\,o_3$	$o_1\,o_4\,o_3\,o_5$	$o_1\,o_4\,o_3\,o_5\,o_2$

Scoring functions

We only have rankings over objects...

$\rightarrow$ *How to compare two allocations ?*

Scoring functions

We only have rankings over objects...

$\rightarrow$ *How to compare two allocations ?*

Two natural assumptions:

1. **Scoring:** We have a common **scoring function** $g : \{1, \ldots, p\} \mapsto \mathbb{N}$ mapping each rank to a utility.
2. **Additivity:** These utilities are **additive**.

Scoring functions

We only have rankings over objects...

$\rightarrow$ *How to compare two allocations ?*

Two natural assumptions:

1. **Scoring:** We have a common **scoring function** $g : \{1, \ldots, p\} \mapsto \mathbb{N}$ mapping each rank to a utility.
2. **Additivity:** These utilities are **additive**.

3 natural scoring functions:

$$\succ_i \quad o_6 \quad o_1 \quad o_4 \quad o_5 \quad o_2 \quad o_3$$

Scoring functions

We only have rankings over objects...

→ *How to compare two allocations ?*

Two natural assumptions:

1. **Scoring:** We have a common **scoring function** $g : \{1, \dots, p\} \mapsto \mathbb{N}$ mapping each rank to a utility.
2. **Additivity:** These utilities are **additive**.

3 natural scoring functions:

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1

Scoring functions

We only have rankings over objects...

$\rightarrow$ *How to compare two allocations ?*

Two natural assumptions:

1. **Scoring:** We have a common **scoring function** $g : \{1, \ldots, p\} \mapsto \mathbb{N}$ mapping each rank to a utility.
2. **Additivity:** These utilities are **additive**.

3 natural scoring functions:

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1
lexicographic	32	16	8	4	2	1

Scoring functions

We only have rankings over objects...

$\rightarrow$ *How to compare two allocations ?*

Two natural assumptions:

1. **Scoring:** We have a common **scoring function** $g : \{1, \ldots, p\} \mapsto \mathbb{N}$ mapping each rank to a utility.
2. **Additivity:** These utilities are **additive**.

3 natural scoring functions:

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1
lexicographic	32	16	8	4	2	1
Quasi-Indifference	$1 + 5\varepsilon$	$1 + 4\varepsilon$	$1 + 3\varepsilon$	$1 + 2\varepsilon$	$1 + \varepsilon$	1

Back to the example

Example

5 objects, 3 agents, $\pi = 12332\ldots$

- $1 : \mathbf{o_1} \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : \mathbf{o_4} \succ \mathbf{o_2} \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ \mathbf{o_3} \succ \mathbf{o_5} \succ o_4 \succ o_2$

With π, agent 1 gets o_1, agent 2 gets $o_2 o_4$, agent 3 gets $o_3 o_5$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332\ldots$

- $1 : \mathbf{o_1} \succ o_2 \succ o_3 \succ o_4 \succ o_5$
- $2 : \mathbf{o_4} \succ \mathbf{o_2} \succ o_5 \succ o_1 \succ o_3$
- $3 : o_1 \succ \mathbf{o_3} \succ \mathbf{o_5} \succ o_4 \succ o_2$

With π, agent 1 gets o_1, agent 2 gets $o_2 o_4$, agent 3 gets $o_3 o_5$

- **Borda:** $u_1(\pi) = 5$; $u_2(\pi) = 5 + 4 = 9$; $u_3(\pi) = 4 + 3 = 7$.
- **lexicographic:** $u_1(\pi) = 16$; $u_2(\pi) = 24$; $u_3(\pi) = 12$.
- **QI:** $u_1(\pi) = 1 + 4\varepsilon$; $u_2(\pi) = 2 + 7\varepsilon$; $u_3(\pi) = 2 + 5\varepsilon$.

Social welfare

We use a **collective utility function** to aggregate the individual utilities.

Two well-known functions:

- **utilitarian:** $F(u_1, \ldots, u_n) = \sum_{i=1,\ldots,n} u_i$;
- **egalitarian:** $F(u_1, \ldots, u_n) = \min_{i=1,\ldots,n} u_i$.

Uncertainty

The procedure is elicitation-free...

$\rightarrow$ *Which information can we use to find the best sequence ?*

Uncertainty

The procedure is elicitation-free...

→ *Which information can we use to find the best sequence ?*

The CA has a **probabilistic model** of the preferences:

- **Full independence** : each profile $R = \langle \succ_1, \dots, \succ_n \rangle$ is equally probable
- **Full correlation** : all the agents have the same ranking ($R = \langle \succ, \dots, \succ \rangle$)

Uncertainty

The procedure is elicitation-free…

→ *Which information can we use to find the best sequence ?*

The CA has a **probabilistic model** of the preferences:

- **Full independence** : each profile $R = \langle \succ_1, \ldots, \succ_n \rangle$ is equally probable
- **Full correlation** : all the agents have the same ranking ($R = \langle \succ, \ldots, \succ \rangle$)

Expected individual and collective utilities:

$$\overline{u(i,\pi)} = \sum_{R \in Prof(\mathcal{N},\mathcal{O})} Pr(R) \times u_i(\pi, R).$$

$$\overline{sw_F(\pi)} = F(\overline{u(1,\pi)}, \ldots, \overline{u(n,\pi)}).$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?
3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3, \pi)} =$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?
3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3, \pi)} = \frac{1}{\binom{5}{2}} \times (3 + 2)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.

What is agent 3's expected utility with this sequence ?

3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + \frac{1}{\binom{5}{2}} \times (4+2)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?
3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + 0.6 + \frac{1}{\binom{5}{2}} \times (5 + 2)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.

What is agent 3's expected utility with this sequence ?

3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + 0.6 + 0.6 + \frac{2}{\binom{5}{2}} \times (4+3)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.

What is agent 3's expected utility with this sequence ?

3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + 0.6 + 0.6 + 1.4 + \frac{2}{\binom{5}{2}} \times (5+3)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?
3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + 0.6 + 0.6 + 1.4 + 1.6 + \frac{\binom{3}{2}}{\binom{5}{2}} \times (5+4)$$

Back to the example

Example

5 objects, 3 agents, $\pi = 12332$, $g = g_{Borda}$, full independence.
What is agent 3's expected utility with this sequence ?
3's preferences: $o_? \succ o_? \succ o_? \succ o_? \succ o_?$

$$\overline{u(3,\pi)} = 0.5 + 0.6 + 0.6 + 1.4 + 1.6 + 2.7 = \mathbf{7.5}$$

Summary

- **Instance:**
  - a number of agents n
  - a number of objects p
  - a scoring function g
  - a correlation profile $Corr \in \{FC, FI\}$
  - a collective utility function F
- **Question:**
  - *What is the policy π maximizing $\overline{sw_F(\pi)}$, under correlation profile Corr ?*

Some general results

1. Full correlation

Some general results

1. Full correlation

Utilitarian CUF (sum)

All policies have the same expected value!

Some general results

1. Full correlation

Utilitarian CUF (sum)

All policies have the same expected value!

Egalitarian CUF (min)

Sequential allocation is **NP**-complete.

(Reduction from [PARTITION])

Some general results

1. Full correlation

Utilitarian CUF (sum)

All policies have the same expected value!

Egalitarian CUF (min)

Sequential allocation is **NP**-complete.

(Reduction from [PARTITION])

What about...

- ...lexicographic scoring ?
- ...quasi-indifference scoring ?
- ...Borda scoring ?

Lexicographic scoring

$\succ_i$	o_6	$\succ$	o_1	$\succ$	o_4	$\succ$	o_5	$\succ$	o_2	$\succ$	o_3
lexicographic	32	$>>$	16	$>>$	8	$>>$	4	$>>$	2	$>>$	1

Egalitarian CUF (min)

Optimal policies: $\sigma(1)\sigma(2)\ldots\sigma(n-1)\sigma(n)^{p-n+1}$ (where σ is a permutation of $\{1,\ldots,n\}$)

Lexicographic scoring

$\succ_i$	o_6	$\succ$	o_1	$\succ$	o_4	$\succ$	o_5	$\succ$	o_2	$\succ$	o_3
lexicographic	32	$>>$	16	$>>$	8	$>>$	4	$>>$	2	$>>$	1

Egalitarian CUF (min)

Optimal policies: $\sigma(1)\sigma(2)\ldots\sigma(n-1)\sigma(n)^{p-n+1}$ (where σ is a permutation of $\{1,\ldots,n\}$)

Example: $\pi = 123333$

- $\overline{u(1,\pi)} = 32$
- $\overline{u(2,\pi)} = 16$
- $\overline{u(3,\pi)} = 8 + 4 + 2 + 1 = 15$

Borda scoring

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1

Borda scoring

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1

This is **polynomial** in p (dynamic programming algorithm).

Borda scoring

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Borda	6	5	4	3	2	1

This is **polynomial** in p (dynamic programming algorithm).

This is **not** polynomial in n.
There is a mistake in the paper!

QI scoring

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Quasi-Indifference	$1+5\varepsilon$	$1+4\varepsilon$	$1+3\varepsilon$	$1+2\varepsilon$	$1+\varepsilon$	1

Egalitarian CUF (min)

Comes down to solving the Borda case!

QI scoring

$\succ_i$	o_6	o_1	o_4	o_5	o_2	o_3
Quasi-Indifference	$1+5\varepsilon$	$1+4\varepsilon$	$1+3\varepsilon$	$1+2\varepsilon$	$1+\varepsilon$	1

Egalitarian CUF (min)

Comes down to solving the Borda case!

Intuition:

- let $m = \lfloor \frac{p}{n} \rfloor$ and $q = p - nm$
- Optimal policies: $\pi = \underbrace{1122}_{n-q \text{ agents}} \overbrace{333444}^{q \text{ agents}}$ and $\pi' = \underbrace{1221}_{n-q \text{ agents}} \overbrace{333444}^{q \text{ agents}}$
- The q last agents are OK $\rightarrow u \geq m+1$
- The $n-q$ first agents: $u = m + x \cdot \varepsilon$ ($x \rightarrow$ Borda)

A complex problem. . .

2. Full independence

A complex problem. . .

2. Full independence

Bad news:

- Even computing the expected utility of a given sequence seems complex (**NP**-hardness conjectured).
- **Intuition:** we have to build a complete search tree to compute it.
- Computing the best sequence is of course even harder!

A complex problem. . .

2. Full independence

Bad news:

- Even computing the expected utility of a given sequence seems complex (**NP**-hardness conjectured).
- **Intuition:** we have to build a complete search tree to compute it.
- Computing the best sequence is of course even harder!

Precise complexity of the problems of:

- computing the expected collective utility of a sequence
- computing the best sequence

are still **open questions**!

A complex problem. . .

Good news:

- Precise complexity of the problems are still open questions
- When $p \to +\infty$, allocations of the form $\sigma_1\sigma_2 \ldots \sigma_k\theta$, where $\sigma_1, \ldots, \sigma_k$, are permutations of $\{1, \ldots, n\}$ (and $p = kn + q$), tend to be optimal.
  Example: 123 123 321 231 132 321 123 . . .

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4		
5		
6		
8		
10		

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4	1221	
5		
6		
8		
10		

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4	1221	1233
5		
6		
8		
10		

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4	1221	1233
5	11222	12332
6		
8		
10		

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4	1221	1233
5	11222	12332
6	121221	123321
8	12212112	11332232
10	1221121221	1231223133

Some examples

Assumptions: Full independence, egalitarian CUF, Borda scoring function.

p	$n = 2$	$n = 3$
4	1221	1233
5	11222	12332
6	121221	123321
8	12212112	11332232
10	1221121221	1231223133

Other examples on
`http://recherche.noiraudes.net/en/sequences.php`

Strategy-proofness ?

Strategy-proofness ?

The cheating agent knows:

- the sequence
- the others' preferences

Strategy-proofness ?

The cheating agent knows:

- the sequence
- the others' preferences

She wants:

- to securely get one target bundle $\mathcal{S}$
- or to get the best bundle she can get.

Strategy-proofness ?

The cheating agent knows:

- the sequence
- the others' preferences

She wants:

- to securely get one target bundle $\mathcal{S}$
- or to get the best bundle she can get.

Her only possible cheating actions:

- choose at given steps **not** to pick her preferred objects.

Strategy-proofness ?

The cheating agent knows:

- the sequence
- the others' preferences

She wants:

- to securely get one target bundle $\mathcal{S}$
- or to get the best bundle she can get.

Her only possible cheating actions:

- choose at given steps **not** to pick her preferred objects.

Some results

- All the agents have the same preferences $\rightarrow$ strategy-proof.
- General case $\rightarrow$ poly-time algorithm to get $\mathcal{S}$ for sure if it is possible.
- Lexico. scoring $\rightarrow$ optimal strategy computable in polynomial time.

Conclusion & Future work

- A simple and intuitive sequential allocation procedure
- Some complexity results for the FC case
- A basic algorithm but no clever result for the FI case
- Some strategical issues studied

Conclusion & Future work

- A simple and intuitive sequential allocation procedure
- Some complexity results for the FC case
- A basic algorithm but no clever result for the FI case
- Some strategical issues studied

- Missing complexity results ?
- Envy-freeness ?
- Game-theoretical analysis ?
- Using multiobjective MDP to model and solve the problem ?