



Picking Sequences for Resource Allocation

Sylvain Bouveret

LIG – Grenoble INP / Ensimag

Jérôme Lang

LAMSADE CNRS – Université Paris Dauphine

Michel Lemaître

Formerly Onera Toulouse

Labex CIMI Workshop on Decision Making

Toulouse, April 26, 2017



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...



Sequences

- Natural and simple protocol
- Used in practice
- **Preference elicitation-free**

- Board games
- Draft mechanisms (sport)
- Course allocation (Harvard Business School)
- ...





What I will present today...

- 1 | **Optimal sequence:** We “feel” that $ABCCBA$ is **fairer** than $AABBCC\dots$
→ *What is the **fairest** sequence ?*
- 2 | **Manipulation:** What if the agents act strategically?
- 3 | **Efficiency:** Picking sequences as optimization procedure.

The model

What is the best sequence?



Agents, objects, preferences...

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A central authority (CA) must find a **policy** (a sequence of agents)
 $\pi : \{1, \dots, p\} \rightarrow \{A, B, \dots, x\}$
- The central authority assumes that each agent i has a (private) **ranking** \succ_i over \mathcal{O} (ex: $6 \succ 1 \succ 4 \succ 5 \succ 2 \succ 3$)



Example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- A : 1 \succ 2 \succ 3 \succ 4 \succ 5
- B : 4 \succ 2 \succ 5 \succ 1 \succ 3
- C : 1 \succ 3 \succ 5 \succ 4 \succ 2



Example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- $A : 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B : 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C : 1 \succ 3 \succ 5 \succ 4 \succ 2$

| | |
|---------------------|-------------|
| k | 0 |
| $s(A)_k^\pi$ | \emptyset |
| $s(B)_k^\pi$ | \emptyset |
| $s(C)_k^\pi$ | \emptyset |
| \mathcal{O}_k^π | \emptyset |



Example

Example

5 objects, 3 agents, $\pi = \mathbf{ABCCB}\dots$

- $A: 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B: 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C: 1 \succ 3 \succ 5 \succ 4 \succ 2$

| | | |
|---------------------|-------------|-------------|
| k | 0 | 1 |
| $s(A)_k^\pi$ | \emptyset | 1 |
| $s(B)_k^\pi$ | \emptyset | \emptyset |
| $s(C)_k^\pi$ | \emptyset | \emptyset |
| \mathcal{O}_k^π | \emptyset | 1 |



Example

Example

5 objects, 3 agents, $\pi = \mathbf{ABCCB}\dots$

- $A: 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B: 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C: 1 \succ 3 \succ 5 \succ 4 \succ 2$

| k | 0 | 1 | 2 |
|---------------------|-------------|-------------|-------------|
| $s(A)_k^\pi$ | \emptyset | 1 | 1 |
| $s(B)_k^\pi$ | \emptyset | \emptyset | 4 |
| $s(C)_k^\pi$ | \emptyset | \emptyset | \emptyset |
| \mathcal{O}_k^π | \emptyset | 1 | 14 |



Example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- A : 1 \succ 2 \succ 3 \succ 4 \succ 5
- B : 4 \succ 2 \succ 5 \succ 1 \succ 3
- C : 1 \succ 3 \succ 5 \succ 4 \succ 2

| k | 0 | 1 | 2 | 3 |
|---------------------|-------------|-------------|-------------|-----|
| $s(A)_k^\pi$ | \emptyset | 1 | 1 | 1 |
| $s(B)_k^\pi$ | \emptyset | \emptyset | 4 | 4 |
| $s(C)_k^\pi$ | \emptyset | \emptyset | \emptyset | 3 |
| \mathcal{O}_k^π | \emptyset | 1 | 14 | 143 |



Example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- $A: 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B: 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C: 1 \succ 3 \succ 5 \succ 4 \succ 2$

| k | 0 | 1 | 2 | 3 | 4 |
|---------------------|-------------|-------------|-------------|-----|------|
| $s(A)_k^\pi$ | \emptyset | 1 | 1 | 1 | 1 |
| $s(B)_k^\pi$ | \emptyset | \emptyset | 4 | 4 | 4 |
| $s(C)_k^\pi$ | \emptyset | \emptyset | \emptyset | 3 | 35 |
| \mathcal{O}_k^π | \emptyset | 1 | 14 | 143 | 1435 |



Example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- $A: 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B: 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C: 1 \succ 3 \succ 5 \succ 4 \succ 2$

| k | 0 | 1 | 2 | 3 | 4 | 5 |
|---------------------|-------------|-------------|-------------|-----|------|-------|
| $s(A)_k^\pi$ | \emptyset | 1 | 1 | 1 | 1 | 1 |
| $s(C)_k^\pi$ | \emptyset | \emptyset | 4 | 4 | 4 | 42 |
| $s(C)_k^\pi$ | \emptyset | \emptyset | \emptyset | 3 | 35 | 35 |
| \mathcal{O}_k^π | \emptyset | 1 | 14 | 143 | 1435 | 14352 |



Scoring functions

We only have rankings over objects...

→ *How to compare two allocations ?*



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**.



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:

| | | | | | | |
|------------|---|---|---|---|---|---|
| γ_i | 6 | 1 | 4 | 5 | 2 | 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 | 6 | 1 | 4 | 5 | 2 | 3 |
|-----------|---|---|---|---|---|---|
| Borda | 6 | 5 | 4 | 3 | 2 | 1 |



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 | 6 | 1 | 4 | 5 | 2 | 3 |
|---------------|----|----|---|---|---|---|
| Borda | 6 | 5 | 4 | 3 | 2 | 1 |
| lexicographic | 32 | 16 | 8 | 4 | 2 | 1 |



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 | 6 | 1 | 4 | 5 | 2 | 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, $\pi = ABCCB\dots$

- $A : 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B : 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C : 1 \succ 3 \succ 5 \succ 4 \succ 2$

With π , agent A gets 1, agent B gets 24, agent C gets 35



Back to the example

Example

5 objects, 3 agents, $\pi = ABCCB\dots$

- $A : 1 \succ 2 \succ 3 \succ 4 \succ 5$
- $B : 4 \succ 2 \succ 5 \succ 1 \succ 3$
- $C : 1 \succ 3 \succ 5 \succ 4 \succ 2$

With π , agent A gets 1, agent B gets 24, agent C gets 35

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



Social welfare

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

Two well-known functions:

- **utilitarian:** $F(u_A, \dots, u_x) = \sum_{i \in \mathcal{N}} u_i$
- **egalitarian:** $F(u_A, \dots, u_x) = \min_{i \in \mathcal{N}} u_i$.



Uncertainty

The procedure is elicitation-free...

→ *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 $R = \langle \succ_A, \dots, \succ_x \rangle$ 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 $R = \langle \succ_A, \dots, \succ_x \rangle$ is equally probable
- **Full correlation** : all the agents have the same ranking ($R = \langle \succ, \dots, \succ \rangle$)

Expected individual and collective utilities:

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

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



Back to the example

Example

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



Back to the example

Example

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

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

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

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



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 | 6 | \succ | 1 | \succ | 4 | \succ | 5 | \succ | 2 | \succ | 3 |
| lexicographic | 32 | \gg | 16 | \gg | 8 | \gg | 4 | \gg | 2 | \gg | 1 |

Egalitarian CUF (min)

Optimal policies: $\sigma(A)\sigma(B)\dots\sigma(x)\sigma(x)^{p-n}$ (where σ is a permutation of $\{A, B, \dots, x\}$)



Lexicographic scoring

| | | | | | | | | | | | |
|---------------|----|----------|----|----------|---|----------|---|----------|---|----------|---|
| γ_i | 6 | γ | 1 | γ | 4 | γ | 5 | γ | 2 | γ | 3 |
| lexicographic | 32 | >> | 16 | >> | 8 | >> | 4 | >> | 2 | >> | 1 |

Egalitarian CUF (min)

Optimal policies: $\sigma(A)\sigma(B)\dots\sigma(x)\sigma(x)^{p-n}$ (where σ is a permutation of $\{A, B, \dots, x\}$)

Example: $\pi = ABCCCC$

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



Borda scoring

| | | | | | | |
|------------|---|---|---|---|---|---|
| γ_i | 6 | 1 | 4 | 5 | 2 | 3 |
| Borda | 6 | 5 | 4 | 3 | 2 | 1 |



Borda scoring

| | | | | | | |
|------------|---|---|---|---|---|---|
| γ_i | 6 | 1 | 4 | 5 | 2 | 3 |
| Borda | 6 | 5 | 4 | 3 | 2 | 1 |

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



QI scoring

| | | | | | | |
|--------------------|--------------------|--------------------|--------------------|--------------------|-------------------|---|
| \succ_i | 6 | 1 | 4 | 5 | 2 | 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

| | | | | | | |
|--------------------|--------------------|--------------------|--------------------|--------------------|-------------------|---|
| \succsim_i | 6 | 1 | 4 | 5 | 2 | 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: $\pi = \underbrace{AABB}_{n-q \text{ agents}} \overbrace{CCDDDD}^{q \text{ agents}}$ and $\pi' = \underbrace{ABBA}_{n-q \text{ agents}} \overbrace{CCDDDD}^{q \text{ agents}}$

- 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

Conjecture (2011)

Computing the expected utility of a sequence is **NP**-complete.
Computing the **optimal** sequence probably harder.



Results

Computing the expected utility of a sequence is **NP**-complete.



Results

Computing the expected utility of a sequence is ~~NP~~-complete
polynomial [Kalinowski et al., 2013].



Kalinowski, T., Narodytska, N., and Walsh, T. (2013).

A social welfare optimal sequential allocation procedure.
In *Proceedings of IJCAI 2013*.



Results

Computing the expected utility of a sequence is ~~NP-complete~~ **polynomial** [Kalinowski et al., 2013].

Computing the **optimal** sequence probably harder.



Kalinowski, T., Narodytska, N., and Walsh, T. (2013).

A social welfare optimal sequential allocation procedure.
In *Proceedings of IJCAI 2013*.



Results

Computing the expected utility of a sequence is **NP-complete polynomial** [Kalinowski et al., 2013].

Computing the **optimal** sequence probably harder.

- the **alternating policy** (*ABABABAB...*) is optimal for Borda, utilitarian social welfare
- complexity unknown for other social welfare and scoring functions (**NP-hardness** conjectured)



Kalinowski, T., Narodytska, N., and Walsh, T. (2013).

A social welfare optimal sequential allocation procedure.
In *Proceedings of IJCAI 2013*.



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 | ABBA | |
| 5 | | |
| 6 | | |
| 8 | | |
| 10 | | |



Some examples

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

| p | $n = 2$ | $n = 3$ |
|-----|---------|---------|
| 4 | ABBA | ABCC |
| 5 | | |
| 6 | | |
| 8 | | |
| 10 | | |



Some examples

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

| p | $n = 2$ | $n = 3$ |
|-----|---------|---------|
| 4 | ABBA | ABCC |
| 5 | AABBB | ABCCB |
| 6 | | |
| 8 | | |
| 10 | | |



Some examples

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

| p | $n = 2$ | $n = 3$ |
|-----|------------|------------|
| 4 | ABBA | ABCC |
| 5 | AABBB | ABCCB |
| 6 | ABABBA | ABCCBA |
| 8 | ABBABAAB | AACCBBCB |
| 10 | ABBAABABBA | ABCABBCACC |



Some examples

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

| p | $n = 2$ | $n = 3$ |
|-----|------------|------------|
| 4 | ABBA | ABCC |
| 5 | AABBB | ABCCB |
| 6 | ABABBA | ABCCBA |
| 8 | ABBABAAB | AACCBBCB |
| 10 | ABBAABABBA | ABCABBCACC |

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

Strategical issues (manipulation)

Is the protocol strategy-proof?



Manipulation?

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A central authority (CA) has chosen a **policy** π and will execute it
- The agents have their own private preferences \rightarrow **picking strategy**.



Manipulation?

Example

2 agents, 4 objects:

- A: $1 \succ 2 \succ 3 \succ 4$
- B: $2 \succ 3 \succ 4 \succ 1$

Sequence $\pi = ABBA \rightarrow \{14|23\}$.



Manipulation?

Example

2 agents, 4 objects:

- A: $1 \succ 2 \succ 3 \succ 4$
- B: $2 \succ 3 \succ 4 \succ 1$

Sequence $\pi = ABBA \rightarrow \{14|23\}$.

What if A knows B's preferences and acts maliciously?



Manipulation?

Example

2 agents, 4 objects:

- A: $1 \succ 2 \succ 3 \succ 4$
- B: $2 \succ 3 \succ 4 \succ 1$

Sequence $\pi = ABBA \rightarrow \{14|23\}$.

What if A knows B's preferences and acts maliciously?

She can manipulate by picking 2 instead of 1 at first step $\rightarrow \{12|34\}$.



More formally

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A **policy** π
- The agents have their own private preferences (which may or may not be additive) and use them for their **picking strategy**.



More formally

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A **policy** π
- The agents have their own private preferences (which may or may not be additive) and use them for their **picking strategy**.

The cheating agent (A) knows:

- the sequence
- her own (general) **picking strategy**
- the others' **picking strategy** (assumed to be **simple** and **deterministic** – as if each agent had an underlying linear order over the objects)



More formally

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A **policy** π
- The agents have their own private preferences (which may or may not be additive) and use them for their **picking strategy**.

The cheating agent (A) knows:

- the sequence
- her own (general) **picking strategy**
- the others' **picking strategy** (assumed to be **simple** and **deterministic** – as if each agent had an underlying linear order over the objects)

She wants:

- to get the best bundle she can get.



More formally

- A set \mathcal{O} of p objects $\{1, \dots, p\}$
- A set \mathcal{N} of n agents $\{A, B, \dots, x\}$
- A **policy** π
- The agents have their own private preferences (which may or may not be additive) and use them for their **picking strategy**.

The cheating agent (A) knows:

- the sequence
- her own (general) **picking strategy**
- the others' **picking strategy** (assumed to be **simple** and **deterministic** – as if each agent had an underlying linear order over the objects)

She wants:

- to get the best bundle she can get.

Her only possible cheating actions:

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



First result

A: *“Can I get S for sure?”*



First result

A: *“Can I get S for sure?”*

Getting a subset for sure

We can answer to that constructively in polynomial time!



First result

A: *“Can I get \mathcal{S} for sure?”*

Getting a subset for sure

We can answer to that constructively in polynomial time!

Idea:

- two agents: pick the objects in \mathcal{S} in reverse order of \succ_B
- more agents:
 - transform agents 2 to $m - 1$ into a single (fake) agent
 - apply the algorithm for 2 agents



General manipulation problem

A: *“What is the best subset I can get?”*



General manipulation problem

A: *“What is the best subset I can get?”*

Idea: Greedily build the optimal achievable subset:

- Find the best object i such that $\{i\}$ is achievable;
- Find the best object j such that $\{i, j\}$ is achievable;
- ...



General manipulation problem

A: *“What is the best subset I can get?”*

Idea: Greedily build the optimal achievable subset:

- Find the best object i such that $\{i\}$ is achievable;
- Find the best object j such that $\{i, j\}$ is achievable;
- ...

Manipulation with additive preferences, two agents

If the manipulator has additive preferences, the optimal manipulation can be computed in polynomial time.



General manipulation problem

A: *“What is the best subset I can get?”*

Idea: Greedily build the optimal achievable subset:

- Find the best object i such that $\{i\}$ is achievable;
- Find the best object j such that $\{i, j\}$ is achievable;
- ...

Manipulation with additive preferences, two agents

If the manipulator has additive preferences, the optimal manipulation can be computed in polynomial time.

Only works **for two agents!**



General manipulation problem

[Aziz et al., 2017]

If the manipulator has additive preferences, the optimal manipulation problem is **NP**-complete.

(reduction from [3-SAT])

- Is there a manipulation that yields a better utility than the truthful report?
 \rightsquigarrow **NP**-complete
- Not true anymore for binary utilities and (ordinal) responsive set extension.



Aziz, H., Bouveret, S., Lang, J., and Mackenzie, S. (2017).

Complexity of manipulating sequential allocation.

In *Proceedings of the 31st AAAI conference on Artificial Intelligence (AAAI'17)*.



Coalitional Manipulation

Example

3 agents, 6 objects:

- A: 1 \succ 2 \succ 5 \succ 4 \succ 3 \succ 6
- B: 1 \succ 3 \succ 5 \succ 2 \succ 4 \succ 6
- C: 2 \succ 3 \succ 4 \succ 1 \succ 5 \succ 6

Sequence $\pi = ABCABC \rightarrow \{15|34|26\}$.



Coalitional Manipulation

Example

3 agents, 6 objects:

- A: 1 \succ 2 \succ 5 \succ 4 \succ 3 \succ 6
- B: 1 \succ 3 \succ 5 \succ 2 \succ 4 \succ 6
- C: 2 \succ 3 \succ 4 \succ 1 \succ 5 \succ 6

Sequence $\pi = ABCABC \rightarrow \{15|34|26\}$.

- If A and B manipulate alone, they cannot do better
- If they cooperate, they can get $\{12|35|46\}$, which is strongly better.



Coalitional Manipulation: Results

Three kinds of manipulation considered here:

- No post-allocation trade allowed between the manipulators
- Post-allocation exchange of goods allowed between the manipulators
- Post-allocation exchange of goods + side-payments allowed



Coalitional Manipulation: Results

Three kinds of manipulation considered here:

- No post-allocation trade allowed between the manipulators
- Post-allocation exchange of goods allowed between the manipulators
- Post-allocation exchange of goods + side-payments allowed

Results:

- No post-allocation trade allowed between the manipulators → NP-complete [PARTITION]
- Post-allocation exchange of goods allowed between the manipulators → NP-complete [PARTITION]
- Post-allocation exchange of goods + side-payments allowed → polynomial (comes down to manipulation by a single agent)



Everyone manipulates...

One manipulator



Everyone manipulates...

One manipulator \rightarrow several manipulators (coalitional manipulation)



Everyone manipulates...

One manipulator → several manipulators (coalitional manipulation)
→ everyone (rational, self-interested) manipulates?



Everyone manipulates...

One manipulator → several manipulators (coalitional manipulation)
→ everyone (rational, self-interested) manipulates?

Game Theory (Subgame Perfect Nash Equilibrium)



Everyone manipulates...

One manipulator \rightarrow several manipulators (coalitional manipulation)
 \rightarrow everyone (rational, self-interested) manipulates?

Game Theory (Subgame Perfect Nash Equilibrium)

- Two agents and additive utilities, precise characterization of the result of every SPNE $((rev(\succ_2), rev(\succ_1), rev(\pi)))$
 [Kalinowski et al., 2013, Kohler and Chandrasekaran, 1971].
- Unbounded number of agents: **PSPACE**-hard [Kalinowski et al., 2013].



Kalinowski, T., Narodytska, N., Walsh, T., and Xia, L. (2013).

Strategic behavior when allocating indivisible goods sequentially.
In Proceedings of AAAI'13.



Kohler, D. A. and Chandrasekaran, R. (1971).

A class of sequential games.
Operations Research, 19(2):270–277.

Picking sequences, fairness, efficiency

A simple yet powerful allocation mechanism



Pareto-efficiency

Pareto-efficiency: we cannot improve the utility of an agent without decreasing the utility of another one.



Pareto-efficiency

Pareto-efficiency: we cannot improve the utility of an agent without decreasing the utility of another one.

Picking sequences: a form of “local efficiency”. At her turn, an agent picks the best item for her.



Pareto-efficiency

Pareto-efficiency: we cannot improve the utility of an agent without decreasing the utility of another one.

Picking sequences: a form of “local efficiency”. At her turn, an agent picks the best item for her.

Is there a link between Pareto-efficiency and picking sequences?



Pareto-efficiency and picking sequences

Yes, there is a link between Pareto-efficiency and picking sequences...

Brams and King (2005)

π Pareto-efficient \Leftrightarrow sequenceable.

(here π Pareto-efficient \Leftrightarrow for all π' , there exists one utility function u compatible with \succ such that π' does not Pareto-dominate π for u .)



Pareto-efficiency and picking sequences

Yes, there is a link between Pareto-efficiency and picking sequences...

Brams and King (2005)

π Pareto-efficient \Leftrightarrow sequenceable.

(here π Pareto-efficient \Leftrightarrow for all π' , there exists one utility function u compatible with \succ such that π' does not Pareto-dominate π for u .)

Possible Pareto-efficiency \leadsto very weak.



Pareto-efficiency and picking sequences

Yes, there is a link between Pareto-efficiency and picking sequences...

Brams and King (2005)

π Pareto-efficient \Leftrightarrow sequenceable.

(here π Pareto-efficient \Leftrightarrow for all π' , there exists one utility function u compatible with \succ such that π' does not Pareto-dominate π for u .)

Possible Pareto-efficiency \leadsto very weak.

However, we have another result:

Proposition [Bouveret and Lemaître, 2016]

Every Pareto-efficient allocation is sequenceable.



Bouveret, S. and Lemaître, M. (2016).

Efficiency and sequenceability in fair division of indivisible goods with additive preferences.

In *Proceedings of the Sixth International Workshop on Computational Social Choice (COMSOC'16)*.



Approximating Envy-freeness

Envy-freeness: everyone thinks her share is better than any other share.



Approximating Envy-freeness

Envy-freeness: everyone thinks her share is better than any other share.

Nice property, but cannot always be satisfied. \rightsquigarrow several relaxations proposed:

- envy-minimization (for several definitions of the quantity of envy);
- envy-freeness up to one good [Budish, 2011]: “I might envy some other agent, but if we remove just one good from the share of this agent, then I do not envy her anymore.”



Budish, E. (2011).

The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes.
Journal of Political Economy, 119(6).



Envy-freeness Up to One Good

[Budish, 2011, Caragiannis et al., 2016]

Any allocation obtained by the round-robin picking sequence is envy-free up to one good.



Budish, E. (2011).

The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes.
Journal of Political Economy, 119(6).



Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016).

The unreasonable fairness of maximum nash welfare.
In *Proceedings of the ACM Conference on Electronic Commerce (EC'16)*.



Price of elicitation-freeness

Another way to ensure fairness and efficiency: maximize a collective utility (social choice) function.



Price of elicitation-freeness

Another way to ensure fairness and efficiency: maximize a collective utility (social choice) function.

Picking sequences: arguably a very simple (and natural) protocol...



Price of elicitation-freeness

Another way to ensure fairness and efficiency: maximize a collective utility (social choice) function.

Picking sequences: arguably a very simple (and natural) protocol...
...but obviously suboptimal



Price of elicitation-freeness

Another way to ensure fairness and efficiency: maximize a collective utility (social choice) function.

Picking sequences: arguably a very simple (and natural) protocol...

...but obviously suboptimal

Question: *What is the loss in social welfare we incur by using such a simple protocol instead of computing the optimal allocation?*

Here we focus on **regular sequences** σ of the kind $(1 \dots n)^*$

(round-robin), but our results are similar for alternating sequences like $(1 \dots nn \dots 1)^*$.



Price of elicitation-freeness

Another way to ensure fairness and efficiency: maximize a collective utility (social choice) function.

Picking sequences: arguably a very simple (and natural) protocol...

...but obviously suboptimal

Question: *What is the loss in social welfare we incur by using such a simple protocol instead of computing the optimal allocation?*

Here we focus on **regular sequences** σ of the kind $(1 \dots n)^*$

(round-robin), but our results are similar for alternating sequences like $(1 \dots nn \dots 1)^*$.

More formally:

- **Multiplicative** Price of Elicitation-Freeness: worst case ratio $u_c^{\text{opt}}/u_c(\sigma)$, for a sequence σ .
- **Additive** Price of Elicitation-Freeness: worst case difference $u_c^{\text{opt}} - u_c(\sigma)$, for a sequence σ .



Some theoretical bounds for Borda

For classical utilitarianism ($\sum_i u_i(\pi)$):

$$1 + \frac{n-1}{m} + \Theta\left(\frac{1}{m^2}\right) \leq MPEF \leq 2 - \frac{1}{n} + \Theta\left(\frac{1}{m^2}\right), \text{ when } m \rightarrow +\infty.$$



Some theoretical bounds for Borda

For **classical utilitarianism** ($\sum_i u_i(\pi)$):

$$1 + \frac{n-1}{m} + \Theta\left(\frac{1}{m^2}\right) \leq MPEF \leq 2 - \frac{1}{n} + \Theta\left(\frac{1}{m^2}\right), \text{ when } m \rightarrow +\infty.$$

For **egalitarianism** ($\min_i u_i(\pi)$):


$$MPEF \leq 2 - \frac{1}{n} + \Theta\left(\frac{1}{m^2}\right), \text{ when } m \rightarrow +\infty.$$

See the paper for more!




Some experimental results for Borda

For classical utilitarianism ($\sum_i u_i(\pi)$):



For egalitarianism ($\min_i u_i(\pi)$):



Conclusion

A take-away message?



Conclusion

- A simple and intuitive sequential allocation procedure (actually already known in sparse literature)
- Finding the best policy?
 - Full Correlation case well understood
 - Full Independence: partial results
- Strategical issues: individual and coalitional manipulation; game-theoretic approach



Very nice properties...

A simple protocol, but with nice features:

- “locally” efficient
- guarantees envy-freeness up to one good
- gives good approximation of social welfare
- also gives good approximation of other fairness properties (e.g. max-min share)