Towards Fairer Collective Decisions

Habilitation Defense

Sylvain Bouveret

LIG, Ensimag-Grenoble INP, Univ. Grenoble-Alpes

January, the 20th, 2025

Towards fairer collective decisions

Collective decision making...

Towards fairer collective decisions

Collective decision making...

- A set of **alternatives** $\mathcal{O}$

Towards fairer collective decisions

Collective decision making...

- A set of **alternatives** $\mathcal{O}$
- A set of **agents** $\mathcal{A} = \{a_1, \dots, a_n\}$...

Towards fairer collective decisions

Collective decision making...

- A set of **alternatives** $\mathcal{O}$
- A set of **agents** $\mathcal{A} = \{a_1, \ldots, a_n\}$...
- ...Expressing **opinions** (preferences) over the alternatives.

Towards fairer collective decisions

Collective decision making...

- A set of **alternatives** $\mathcal{O}$
- A set of **agents** $\mathcal{A} = \{a_1, \ldots, a_n\}$...
- ...Expressing **opinions** (preferences) over the alternatives.

$$\Downarrow$$

Collective opinion, choice of an alternative...

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

Candidate 1

Candidate 2

Candidate 3

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

Candidate 1 Candidate 2 Candidate 3

Voter 1 Voter 2 Voter 3 Voter 4 Voter 5

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

Candidate 1 Candidate 2 Candidate 3

Voter 1 Voter 2 Voter 3 Voter 4 Voter 5

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

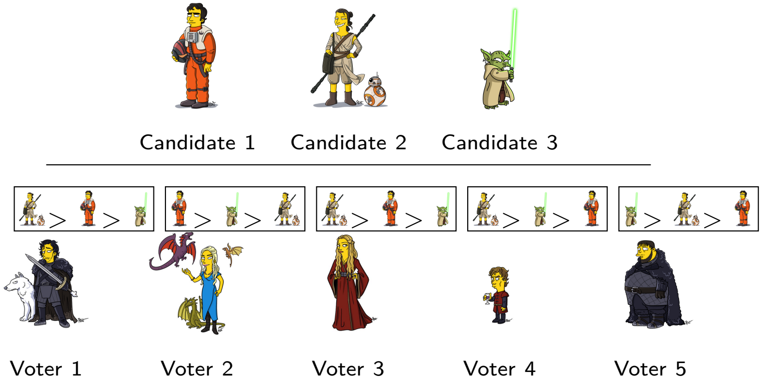

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

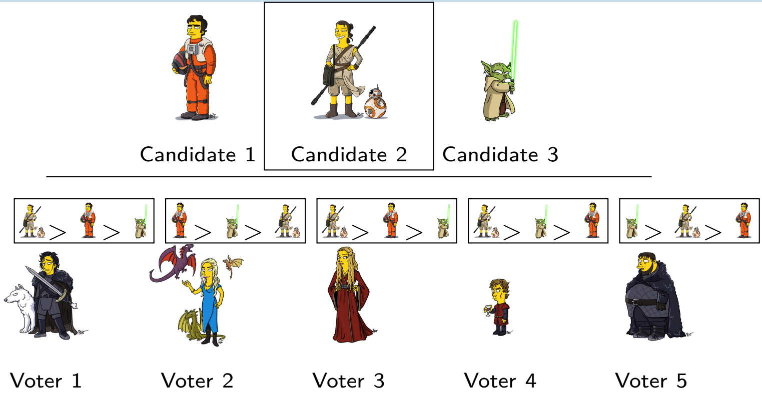

Voting

Problem #1: Voting

We have to elect a representative from a set of m candidates on which the n voters have diverse preferences.

- **Alternatives**: candidates
- **Agents**: voters
- **Preferences**: ballots (linear orders, single-name ballots...)

Applications: political elections, middle or low-stake elections (*e.g* hire a new colleague), choose a restaurant...

Fair division of indivisible goods

Problem #2: Discrete fair division

We have to allocate a set of m indivisible items to n agents having different evaluations of these objects.

Fair division of indivisible goods

Problem #2: Discrete fair division

We have to allocate a set of m indivisible items to n agents having different evaluations of these objects.

Agent 1

Agent 2

Agent 3

Agent 4

Fair division of indivisible goods

Problem #2: Discrete fair division

We have to allocate a set of m indivisible items to n agents having different evaluations of these objects.

Agent 1

Agent 2

Agent 3

Agent 4

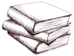

Object 1

Object 2

Object 3

Object 4

Object 5

Fair division of indivisible goods

Problem #2: Discrete fair division

We have to allocate a set of m indivisible items to n agents having different evaluations of these objects.

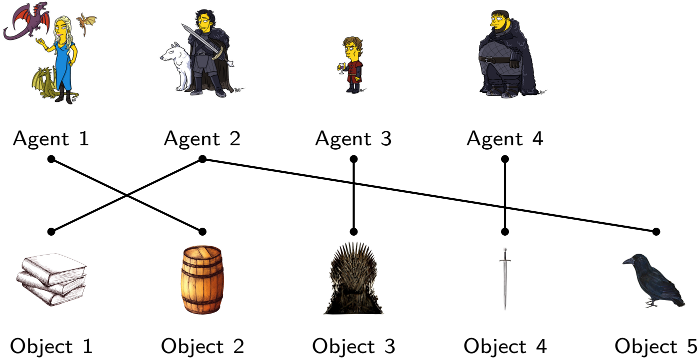

Fair division of indivisible goods

Problem #2: Discrete fair division

We have to allocate a set of m indivisible items to n agents having different evaluations of these objects.

- **Alternatives**: possible allocations (n^m)
- **Agents**: objects consumers (n)
- **Preferences**: utility functions / orders...

Applications: dividing inheritance, allocating lab works to students, papers to reviewers, tasks to robots or machines, tasks in crowdsourcing systems,...

Objectives of the talk

A central topic in these problems: **fairness**...

Objectives of the talk

A central topic in these problems: **fairness**...

How can fairness be formally defined, and how does the use of different fairness notions impact the collective decision and its computation in practice?

Objectives of the talk

A central topic in these problems: **fairness**...

> How can fairness be formally defined, and how does the use of different fairness notions impact the collective decision and its computation in practice?

In this talk:

- Some of the topics I have been working on mostly **between 2011 and 2019**
- All these topics belong to the domain of **Computational Social Choice** (COMSOC) $\approx$ Social Choice Theory $\cap$ Computer Science

Outline

1. Fair enough: fairness beyond proportionality and envy-freeness

2. The unreasonable fairness of picking sequences

3. And the winner is... Alternative (fairer?) voting rules

Fair divison

Fair enough: fairness beyond proportionality and envy-freeness

The fair division problem

You have:

- a finite set of **objects** $\mathcal{O} = \{o_1, \ldots, o_m\}$
- a finite set of **agents** $\mathcal{A} = \{a_1, \ldots, a_n\}$ having some **preferences** on the set of objects they may receive

The fair division problem

You have:

- a finite set of **objects** $\mathcal{O} = \{o_1, \ldots, o_m\}$
- a finite set of **agents** $\mathcal{A} = \{a_1, \ldots, a_n\}$ having some **preferences** on the set of objects they may receive

You want:

- an allocation $\overrightarrow{\pi} : \mathcal{A} \to 2^{\mathcal{O}}$, such that
  - $\pi_i \cap \pi_j = \emptyset$ if $i \neq j$ (preemption),
  - $\bigcup_{a \in \mathcal{A}} \pi_i = \mathcal{O}$ (no free-disposal),
- and which takes into account the agents preferences

The fair division problem

You have:

- a finite set of **objects** $\mathcal{O} = \{o_1, \ldots, o_m\}$
- a finite set of **agents** $\mathcal{A} = \{a_1, \ldots, a_n\}$ having some **preferences** on the set of objects they may receive

You want:

- an allocation $\overrightarrow{\pi} : \mathcal{A} \to 2^{\mathcal{O}}$, such that
  - $\pi_i \cap \pi_j = \emptyset$ if $i \neq j$ (preemption),
  - $\bigcup_{a \in \mathcal{A}} \pi_i = \mathcal{O}$ (no free-disposal),
- and which takes into account the agents preferences

Agents preferences...

1. How to express them formally?
2. How to take them into account to compute an allocation?

Additive fair division

1. **Preferences** – a standard model: **additive preferences**

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$

Additive fair division

1. **Preferences** – a standard model: **additive preferences**

- Ask each agent a_i to give a score $w_i(o)$ to each object o
- If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
- **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
   - **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)
2. **Computing an allocation** – two standard approaches:

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
   - **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)
2. **Computing an allocation** – two standard approaches:

1. Find an allocation $\overrightarrow{\pi}$ that maximizes a **collective utility function**, *e.g.* $uc(\overrightarrow{\pi}) = \min_{a_i \in \mathcal{A}} u_i(\pi_i)$ – egalitarian solution

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
   - **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)
2. **Computing an allocation** – two standard approaches:
   1. Find an allocation $\overrightarrow{\pi}$ that maximizes a **collective utility function**, *e.g.* $uc(\overrightarrow{\pi}) = \min_{a_i \in \mathcal{A}} u_i(\pi_i)$ – egalitarian solution
   2. Find an allocation $\overrightarrow{\pi}$ that satisfies a given **fairness criterion**

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
   - **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)
2. **Computing an allocation** – two standard approaches:

1. Find an allocation $\overrightarrow{\pi}$ that maximizes a **collective utility function**, *e.g.* $uc(\overrightarrow{\pi}) = \min_{a_i \in \mathcal{A}} u_i(\pi_i)$ – egalitarian solution
2. Find an allocation $\overrightarrow{\pi}$ that satisfies a given **fairness criterion**

Additive fair division

1. **Preferences** – a standard model: **additive preferences**
   - Ask each agent a_i to give a score $w_i(o)$ to each object o
   - If a_i receives bundle π, she derives utility $u_i(\pi) = \sum_{o \in \pi} w_i(o)$
   - **Additive preferences**: good compromise between simplicity and expressive power (other models exist: linear orders, k-additive preferences...)

2. **Computing an allocation** – two standard approaches:

1. Find an allocation $\overrightarrow{\pi}$ that maximizes a **collective utility function**, *e.g.* $uc(\overrightarrow{\pi}) = \min_{a_i \in \mathcal{A}} u_i(\pi_i)$ – egalitarian solution
2. Find an allocation $\overrightarrow{\pi}$ that satisfies a given **fairness criterion**

In this part, we focus on the *2nd approach* and investigate how fairness can be formally modeled

Two standard criteria

Envy-freeness (EF) [Foley, 1967]

An allocation $\overrightarrow{\pi}$ is **envy-free** if no agent envies another one, that is, $\forall a_i, a_j,\ u_i(\pi_i) \geq u_i(\pi_j)$.

Two standard criteria

Envy-freeness (EF) [Foley, 1967]

An allocation $\overrightarrow{\pi}$ is **envy-free** if no agent envies another one, that is, $\forall a_i, a_j,\ u_i(\pi_i) \geq u_i(\pi_j)$.

Proportional share (PROP) [Steinhaus, 1948]

An allocation $\overrightarrow{\pi}$ satisfies **proportionality** if every agent gets at least $1/n^{\text{th}}$ of the total value of the objects, that is, $\forall a_i,\ u_i(\pi_i) \geq u_i(\mathcal{O})/n$.

Two standard criteria

Envy-freeness (EF) [Foley, 1967]

An allocation $\overrightarrow{\pi}$ is **envy-free** if no agent envies another one, that is, $\forall a_i, a_j,\ u_i(\pi_i) \geq u_i(\pi_j)$.

Proportional share (PROP) [Steinhaus, 1948]

An allocation $\overrightarrow{\pi}$ satisfies **proportionality** if every agent gets at least $1/n^{\text{th}}$ of the total value of the objects, that is, $\forall a_i,\ u_i(\pi_i) \geq u_i(\mathcal{O})/n$.

Known facts:

- $\overrightarrow{\pi}$ is EF $\Rightarrow$ $\overrightarrow{\pi}$ satisfies PROP
- An envy-free (resp. proportional) allocation may not exist
- Deciding whether an allocation is EF (resp. PROP) is polynomial
- Deciding whether an instance has an EF (resp. PROP) allocation is **NP**-complete [Lipton et al., 2004]

Beyond EF and proportionality

Envy-free or proportional allocations are nice, but...

- (...they can be hard to compute)
- ...they do not always exist (what can we do if there are none?)
- ...there can be potentially many of them (how to choose between them?)

Beyond EF and proportionality

Envy-free or proportional allocations are nice, but...

- (...they can be hard to compute)
- ...they do not always exist (what can we do if there are none?)
- ...there can be potentially many of them (how to choose between them?)

Can we enrich the landscape of fairness properties to overcome these problems?

Max-min share

Let us start with proportionality... A nice property, but sometimes too demanding for indivisible goods.

Max-min share

Let us start with proportionality... A nice property, but sometimes too demanding for indivisible goods.

Example

2 agents, 5 objects, with $w_i(o_i) = 1$ for all a_i and o_j.

Max-min share

Let us start with proportionality... A nice property, but sometimes too demanding for indivisible goods.

Example

2 agents, 5 objects, with $w_i(o_i) = 1$ for all a_i and o_j.

Idea [Budish, 2011]: run a "I cut, you choose" game...

- **divisible** (cake-cutting) setting: the agent obtains a **proportional** share
- **indivisible** setting: yields a weaker guarantee, **max-min share**

Max-min share

Let us start with proportionality... A nice property, but sometimes too demanding for indivisible goods.

Example

2 agents, 5 objects, with $w_i(o_i) = 1$ for all a_i and o_j.

Idea [Budish, 2011]: run a "I cut, you choose" game...

- **divisible** (cake-cutting) setting: the agent obtains a **proportional** share
- **indivisible** setting: yields a weaker guarantee, **max-min share**

Max-min share (MmS)

An allocation $\vec{\pi}$ satisfies **max-min share** if
$\forall a_i,\ u_i(\pi_i) \geq \max_{\vec{\pi}} \min_{a_j \in \mathcal{A}} u_i(\pi_j)$.

Max-min share: known facts

- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]

Max-min share: known facts

- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)

Max-min share: known facts

- $\overrightarrow{\pi}$ satisfies PROP $\Rightarrow$ $\overrightarrow{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]

Max-min share: known facts

- $\overrightarrow{\pi}$ satisfies PROP $\Rightarrow$ $\overrightarrow{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]
- Since then...

Max-min share: known facts

- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]
- Since then...
  - A lot of follow-up works on this question

Max-min share: known facts

- $\overrightarrow{\pi}$ satisfies PROP $\Rightarrow$ $\overrightarrow{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]
- Since then...
  - A lot of follow-up works on this question
  - Complexity of deciding whether there exists a max-min share allocation: still open

Max-min share: known facts

- $\overrightarrow{\pi}$ satisfies PROP $\Rightarrow$ $\overrightarrow{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]
- Since then...
  - A lot of follow-up works on this question
  - Complexity of deciding whether there exists a max-min share allocation: still open
  - The best approximation factor so far is $\frac{3}{4} + \frac{3}{3836}$ [Akrami and Garg, 2024]

Max-min share: known facts

- $\overrightarrow{\pi}$ satisfies PROP $\Rightarrow$ $\overrightarrow{\pi}$ satisfies MmS [B. and Lemaître, AAMAS'14]
- Conjecture [B. and Lemaître, AAMAS'14]: a max-min share allocation always exists (proved for special cases, no counterexample found on thousands of random instances)
- Actually proved to be wrong with a very tricky counterexample [Procaccia and Wang, 2014]
- Since then...
  - A lot of follow-up works on this question
  - Complexity of deciding whether there exists a max-min share allocation: still open
  - The best approximation factor so far is $\frac{3}{4} + \frac{3}{3836}$ [Akrami and Garg, 2024]
  - In practice, a MmS allocation exists with very high probability [Kurokawa et al., 2016, Amanatidis et al., 2017]

Min-max share

- Max-min share: *"I cut, you choose (I choose last)"*

Min-max share

- Max-min share: *"I cut, you choose (I choose last)"*
- **Idea** [B. and Lemaître, AAMAS'14]: why not do the opposite (*"Someone cuts, I choose first"*) ?

→ Min-max share

Min-max share

- Max-min share: *"I cut, you choose (I choose last)"*
- **Idea** [B. and Lemaître, AAMAS'14]: why not do the opposite (*"Someone cuts, I choose first"*) ?

→ Min-max share

Min-max share (mMs)

An allocation $\vec{\pi}$ satisfies **min-max share** if

$\forall a_i,\ u_i(\pi_i) \geq \min_{\vec{\pi}} \max_{a_j \in \mathcal{A}} u_i(\pi_j)$.

Min-max share

- Max-min share: *"I cut, you choose (I choose last)"*
- **Idea** [B. and Lemaître, AAMAS'14]: why not do the opposite (*"Someone cuts, I choose first"*) ?

$\rightarrow$ Min-max share

Min-max share (mMs)

An allocation $\vec{\pi}$ satisfies **min-max share** if
$\forall a_i,\ u_i(\pi_i) \geq \min_{\vec{\pi}} \max_{a_j \in \mathcal{A}} u_i(\pi_j)$.

Interestingly,

- $\vec{\pi}$ satisfies mMS $\Rightarrow$ $\vec{\pi}$ satisfies PROP
- $\vec{\pi}$ is EF $\Rightarrow$ $\vec{\pi}$ satisfies mMS

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money
  - Fix a price to each item

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money
  - Fix a price to each item
  - Let the agents buy whatever they want

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money
  - Fix a price to each item
  - Let the agents buy whatever they want
  - If the market clears, the allocation is **fair** and **efficient**

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money
  - Fix a price to each item
  - Let the agents buy whatever they want
  - If the market clears, the allocation is **fair** and **efficient**
- Deciding whether there exists a CEEI is **NP**-hard [Brânzei et al., 2015] but we can compute one (if there is one) using linear programming [B. and Lemaître, COMSOC'16]

CEEI

Competitive Equilibrium from Equal Incomes (CEEI):

- Standard notion in economics, subcase of the Fisher model
- Introduced recently in computer science
- Idea:
  - Endow each agent with the same amount of money
  - Fix a price to each item
  - Let the agents buy whatever they want
  - If the market clears, the allocation is **fair** and **efficient**
- Deciding whether there exists a CEEI is **NP**-hard [Brânzei et al., 2015] but we can compute one (if there is one) using linear programming [B. and Lemaître, COMSOC'16]
- $\overrightarrow{\pi}$ is a CEEI $\Rightarrow$ $\overrightarrow{\pi}$ is EF

A scale of criteria

Let us wrap things up...

A scale of criteria

Let us wrap things up...

- $\vec{\pi}$ is a CEEI $\Rightarrow$ $\vec{\pi}$ is EF
- $\vec{\pi}$ is EF $\Rightarrow$ $\vec{\pi}$ satisfies mMS
- $\vec{\pi}$ satisfies mMS $\Rightarrow$ $\vec{\pi}$ satisfies PROP
- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS
- MmS almost always satisfiable

A scale of criteria

Let us wrap things up...

- $\vec{\pi}$ is a CEEI $\Rightarrow$ $\vec{\pi}$ is EF
- $\vec{\pi}$ is EF $\Rightarrow$ $\vec{\pi}$ satisfies mMS
- $\vec{\pi}$ satisfies mMS $\Rightarrow$ $\vec{\pi}$ satisfies PROP
- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS
- MmS almost always satisfiable

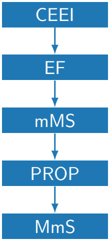

A scale of criteria

Let us wrap things up...

- $\vec{\pi}$ is a CEEI $\Rightarrow$ $\vec{\pi}$ is EF
- $\vec{\pi}$ is EF $\Rightarrow$ $\vec{\pi}$ satisfies mMS
- $\vec{\pi}$ satisfies mMS $\Rightarrow$ $\vec{\pi}$ satisfies PROP
- $\vec{\pi}$ satisfies PROP $\Rightarrow$ $\vec{\pi}$ satisfies MmS
- MmS almost always satisfiable

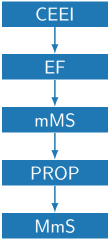

Our approach to fairness [B. and Lemaître, JAAMAS'15]:

1. Determine the highest satisfiable criterion
2. Find an allocation that satisfies this criterion
3. Explain to the upset agents that we cannot do much better

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

- ...start with envy-freeness,
- then propose a relaxation that can be more easily satisfied

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

- ...start with envy-freeness,
- then propose a relaxation that can be more easily satisfied

A standard relaxation: measure of envy [Lipton et al., 2004]

- Pairwise envy: $pe(a_i, a_j, \overrightarrow{\pi}) = \max\{0, u_i(\pi_j) - u_i(\pi_i)\}$
- Individual envy: $e(a_i, \overrightarrow{\pi}) = \max_{a_j \in \mathcal{A}}(pe(a_i, a_j, \overrightarrow{\pi}))$

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

- ...start with envy-freeness,
- then propose a relaxation that can be more easily satisfied

A standard relaxation: measure of envy [Lipton et al., 2004]

- Pairwise envy: $pe(a_i, a_j, \overrightarrow{\pi}) = \max\{0, u_i(\pi_j) - u_i(\pi_i)\}$
- Individual envy: $e(a_i, \overrightarrow{\pi}) = \max_{a_j \in \mathcal{A}}(pe(a_i, a_j, \overrightarrow{\pi}))$
- Collective envy:

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

- ...start with envy-freeness,
- then propose a relaxation that can be more easily satisfied

A standard relaxation: measure of envy [Lipton et al., 2004]

- Pairwise envy: $pe(a_i, a_j, \overrightarrow{\pi}) = \max\{0, u_i(\pi_j) - u_i(\pi_i)\}$
- Individual envy: $e(a_i, \overrightarrow{\pi}) = \max_{a_j \in \mathcal{A}}(pe(a_i, a_j, \overrightarrow{\pi}))$
- Collective envy:
    - sum of individual envies [Lipton et al., 2004]

Relaxing envy-freeness

Our scale of fairness is made of properties conveying different rationales (and rather unexpectedly related to form a scale)

Another approach is possible...

- ...start with envy-freeness,
- then propose a relaxation that can be more easily satisfied

A standard relaxation: measure of envy [Lipton et al., 2004]

- Pairwise envy: $pe(a_i, a_j, \overrightarrow{\pi}) = \max\{0, u_i(\pi_j) - u_i(\pi_i)\}$
- Individual envy: $e(a_i, \overrightarrow{\pi}) = \max_{a_j \in \mathcal{A}}(pe(a_i, a_j, \overrightarrow{\pi}))$
- Collective envy:
  - sum of individual envies [Lipton et al., 2004]
  - a balanced approach, like OWA [Shams, Beynier, B. and Maudet, ADT'21]

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$
- An EF1 allocation **always exists** (and easy to compute)

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?
- Obviously, $(\vec{\pi} \models \text{EF}) \Rightarrow (\vec{\pi} \models \text{EF1})$
- An EF1 allocation **always exists** (and easy to compute)
- A variation of EF1, **envy-free up to any good** (EFX) [Caragiannis et al., 2016]: a_i envies a_j? Is it still the case if we remove **any** item from π_j?

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$
- An EF1 allocation **always exists** (and easy to compute)

- A variation of EF1, **envy-free up to any good** (EFX) [Caragiannis et al., 2016]: a_i envies a_j? Is it still the case if we remove **any** item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EFX}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$

Envy-free up to one good

- **Rationale behind the measure of envy:** if an EF allocation does not exist, agents should be ready to accept a small amount of envy
- Same idea behind **envy-free up to one good** (EF1) [Budish, 2011]: a_i envies a_j? Is it still the case if we remove one (the highest ranked) item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$
- An EF1 allocation **always exists** (and easy to compute)

- A variation of EF1, **envy-free up to any good** (EFX) [Caragiannis et al., 2016]: a_i envies a_j? Is it still the case if we remove **any** item from π_j?
- Obviously, $(\overrightarrow{\pi} \models \text{EF}) \Rightarrow (\overrightarrow{\pi} \models \text{EFX}) \Rightarrow (\overrightarrow{\pi} \models \text{EF1})$
- Complexity of deciding whether there exists an EFX allocation: still open

Landscape, completed

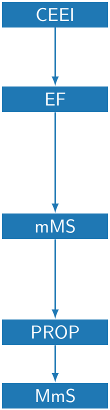

Landscape, completed

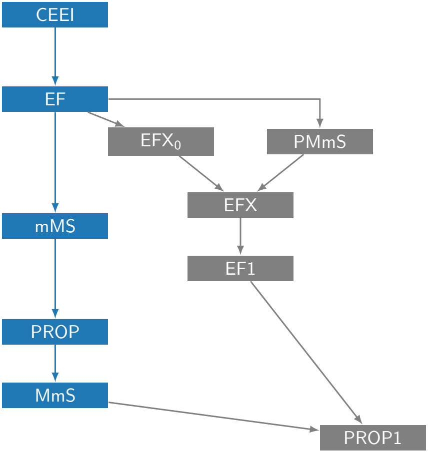

Epistemic envy

Another relaxation of EF...

- EF assumes that the agents have **full knowledge** of the other shares
- In practice, this is unrealistic
- If we assume that they only know their own share → **epistemic envy-freeness** [Aziz, B., Caragiannis, Giagkousi and Lang, AAAI'18]

Epistemic envy

Another relaxation of EF...

- EF assumes that the agents have **full knowledge** of the other shares
- In practice, this is unrealistic
- If we assume that they only know their own share → **epistemic envy-freeness** [Aziz, B., Caragiannis, Giagkousi and Lang, AAAI'18]

Epistemic envy-freeness (EEF)

An agent a_i is EEF in $\overrightarrow{\pi}$ if there is an alternative allocation $\overrightarrow{\pi}'$ such that $\pi'_i = \pi_i$ and a_i is EF in π'.

Epistemic envy

Another relaxation of EF...

- EF assumes that the agents have **full knowledge** of the other shares
- In practice, this is unrealistic
- If we assume that they only know their own share $\rightarrow$ **epistemic envy-freeness** [Aziz, B., Caragiannis, Giagkousi and Lang, AAAI'18]

Epistemic envy-freeness (EEF)

An agent a_i is EEF in $\overrightarrow{\pi}$ if there is an alternative allocation $\overrightarrow{\pi}'$ such that $\pi'_i = \pi_i$ and a_i is EF in π'.

- $(\overrightarrow{\pi} \models EF) \Rightarrow (\overrightarrow{\pi} \models EEF) \Rightarrow (\overrightarrow{\pi} \models mMS)$
- Extended recently to EEFX [Caragiannis et al., 2023] and related concepts
- Intermediate concept: the agents know **some** agents, *via* a social graph G $\rightarrow$ G-EEF

Landscape, completed

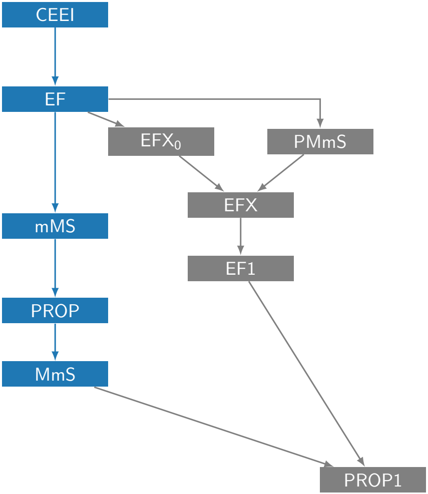

Landscape, completed

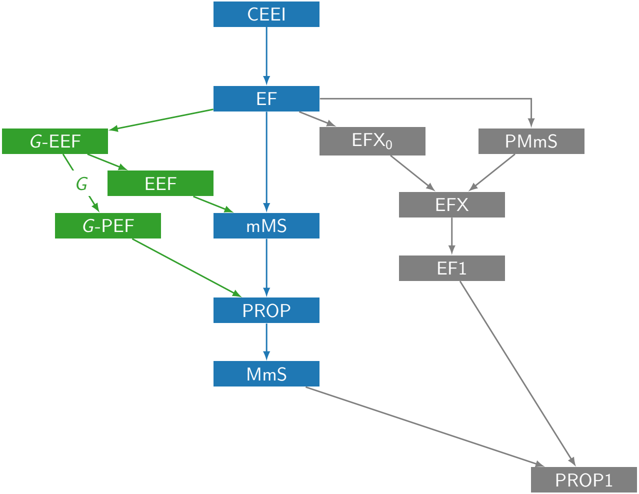

Landscape, completed

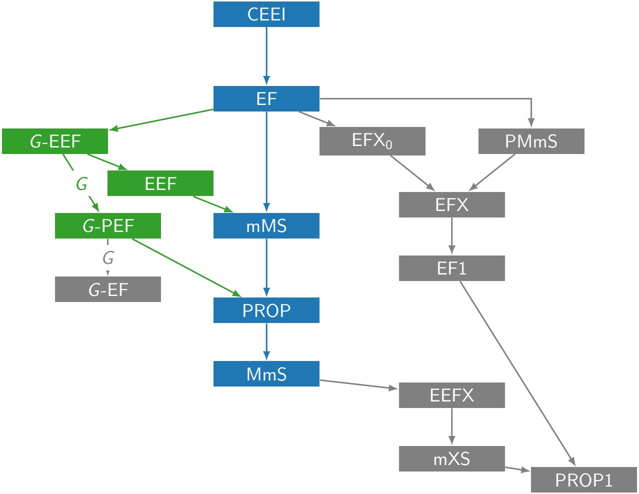

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion
- Suppose a_i envies a_j but noone agrees with a_i: a_i's envy towards a_j is **socially unsupported** [Parijs, 1997]

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion
- Suppose a_i envies a_j but noone agrees with a_i: a_i's envy towards a_j is **socially unsupported** [Parijs, 1997]
- If only some agents support this envy → K-**approval envy** [Shams, Beynier, B. and Maudet, JAIR'22]

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion
- Suppose a_i envies a_j but noone agrees with a_i: a_i's envy towards a_j is **socially unsupported** [Parijs, 1997]
- If only some agents support this envy $\rightarrow$ K-**approval envy** [Shams, Beynier, B. and Maudet, JAIR'22]

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion
- Suppose a_i envies a_j but noone agrees with a_i: a_i's envy towards a_j is **socially unsupported** [Parijs, 1997]
- If only some agents support this envy $\rightarrow$ K-**approval envy** [Shams, Beynier, B. and Maudet, JAIR'22]

K-approval envy (K-app envy)

a_i K-app envies a_j if $\exists$ a subset $\mathcal{A}_K$ of K agents including a_i such that $\forall a_k \in \mathcal{A}_K$, $u_k(\pi_i) < u_k(\pi_j)$

Envy approved by the society

- **Epistemic envy-freeness**: envy is a **knowledge-sensitive** notion
- Another approach: envy as a **subjective** notion
- Suppose a_i envies a_j but noone agrees with a_i: a_i's envy towards a_j is **socially unsupported** [Parijs, 1997]
- If only some agents support this envy $\rightarrow$ K-**approval envy** [Shams, Beynier, B. and Maudet, JAIR'22]

K-approval envy (K-app envy)

a_i K-app envies a_j if $\exists$ a subset $\mathcal{A}_K$ of K agents including a_i such that $\forall a_k \in \mathcal{A}_K,\ u_k(\pi_i) < u_k(\pi_j)$

- $\overrightarrow{\pi}$ is (K-app envy)-free $\Rightarrow$ $\overrightarrow{\pi}$ is (($K+1$)-app envy)-free
- Finding the minimum K so that $\overrightarrow{\pi}$ is (K-app envy)-free is **NP**-complete
- We can extend this concept to K-**app non-proportionality**

Landscape, completed

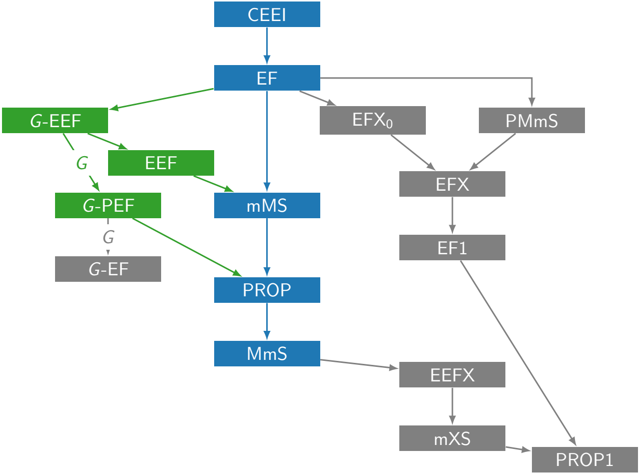

Landscape, completed

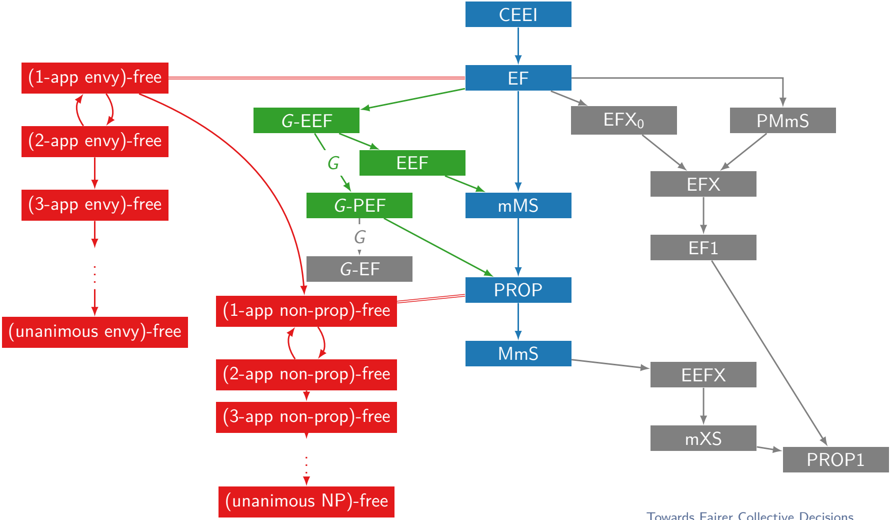

Fair division

The unreasonable fairness of picking sequences

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.

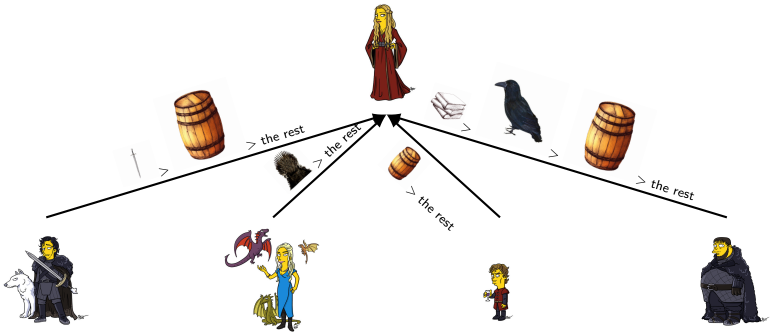

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.

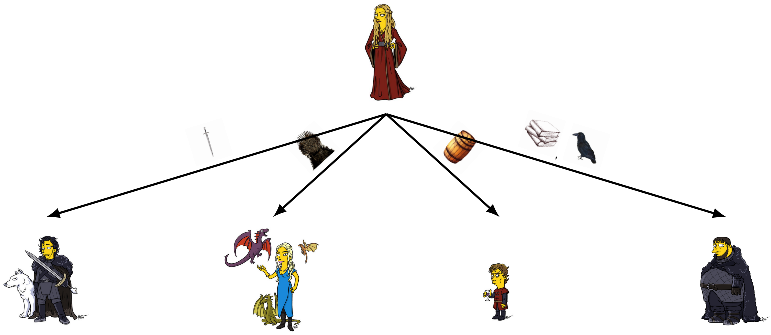

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.

- Needs computational power
- Should be trusted

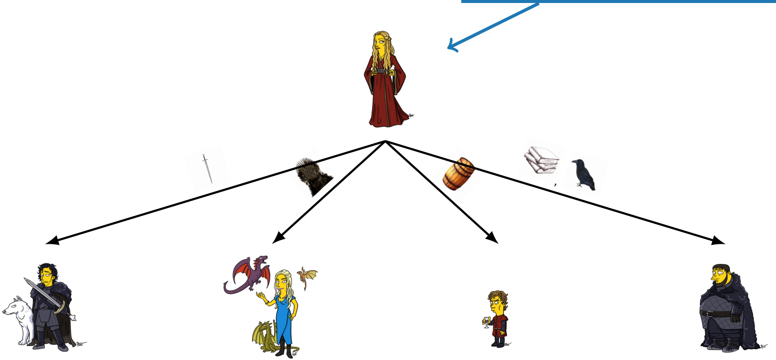

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.

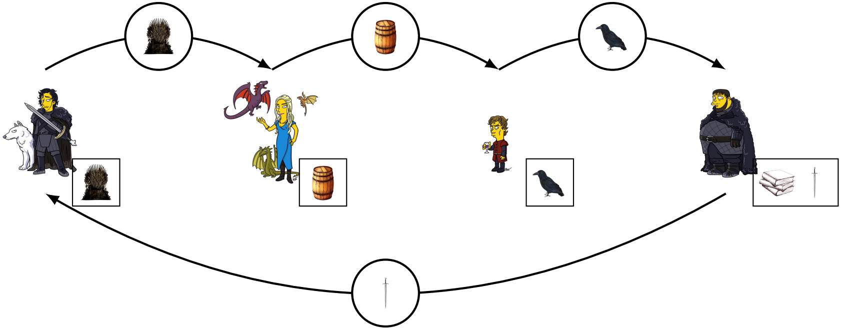

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

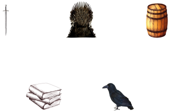

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

How to compute a fair division...

1. So far, what we have done: (i) ask the agents to give their preferences, then (ii) use a (centralized) collective decision making procedure.
2. Start from a random allocation and ask the agents to negotiate.
3. Use an interactive protocol like **picking sequences**.

In this part, we will focus on *picking sequences* (but also talk a little bit about negotiation)

Picking sequences

Picking sequences are...

- natural and simple
- used in practice (board games, draft mechanisms, course allocation...)
- preference elicitation-free

Picking sequences

Picking sequences are...

- natural and simple
- used in practice (board games, draft mechanisms, course allocation...)
- preference elicitation-free

Is this protocol compatible with **fairness requirements**?

Picking sequences

Picking sequences are...

- natural and simple
- used in practice (board games, draft mechanisms, course allocation...)
- preference elicitation-free

Is this protocol compatible with **fairness requirements**?

It depends...

Picking sequences

Picking sequences are...

- natural and simple
- used in practice (board games, draft mechanisms, course allocation...)
- preference elicitation-free

Is this protocol compatible with **fairness requirements**?

It depends... For instance, we "feel" that *ABCCBA* is **fairer** than *AABBCC*

Picking sequences

Picking sequences are...

- natural and simple
- used in practice (board games, draft mechanisms, course allocation...)
- preference elicitation-free

Is this protocol compatible with **fairness requirements**?

It depends... For instance, we "feel" that *ABCCBA* is **fairer** than *AABBCC*

Question

What is the *fairest* sequence?

The formal model

More precisely...

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents $(A, B, C, \ldots)$ and a set $\mathcal{O}$ of m objects

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents $(A, B, C, ...)$ and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \ldots, m\} \rightarrow \mathcal{A}$

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents ($A, B, C, \ldots$) and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \ldots, m\} \rightarrow \mathcal{A}$

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents ($A, B, C, \ldots$) and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \ldots, m\} \to \mathcal{A}$

Fair policy? We **don't know** the agents' preferences...

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents $(A, B, C, ...)$ and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \dots, m\} \rightarrow \mathcal{A}$

Fair policy? We **don't know** the agents' preferences...

Our proposal [B. and Lang, IJCAI'11]: find a policy maximizing **expected social welfare** given some assumptions on the preferences:

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents $(A, B, C, ...)$ and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \dots, m\} \to \mathcal{A}$

Fair policy? We **don't know** the agents' preferences...

Our proposal [B. and Lang, IJCAI'11]: find a policy maximizing **expected social welfare** given some assumptions on the preferences:

1. each agent a_i has a (private) **ranking** $\succ_i$ over $\mathcal{A}$ drawn from a prior probability distribution Ψ (Full Independence (FI) or full correlation (FC))

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents $(A, B, C, \dots)$ and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \dots, m\} \to \mathcal{A}$

Fair policy? We **don't know** the agents' preferences...

Our proposal [B. and Lang, IJCAI'11]: find a policy maximizing **expected social welfare** given some assumptions on the preferences:

1. each agent a_i has a (private) **ranking** $\succ_i$ over $\mathcal{A}$ drawn from a prior probability distribution Ψ (Full Independence (FI) or full correlation (FC))
2. ...rankings are lifted to utilities using a **scoring function** g, *e.g* Borda, lexicographic, quasi-indifference (QI)

The formal model

More precisely...

- **We have**: A set $\mathcal{A}$ of n agents ($A, B, C, \ldots$) and a set $\mathcal{O}$ of m objects
- **We want**: a **fair policy** (a sequence of agents) $\sigma : \{1, \ldots, m\} \to \mathcal{A}$

Fair policy? We **don't know** the agents' preferences...

Our proposal [B. and Lang, IJCAI'11]: find a policy maximizing **expected social welfare** given some assumptions on the preferences:

1. each agent a_i has a (private) **ranking** $\succ_i$ over $\mathcal{A}$ drawn from a prior probability distribution Ψ (Full Independence (FI) or full correlation (FC))
2. ...rankings are lifted to utilities using a **scoring function** g, *e.g* Borda, lexicographic, quasi-indifference (QI)
3. ...individual utilities are aggregated to collective utilities using a **social welfare function** sw, *e.g* egalitarian (min) or utilitarian (sum)

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Examples (Full independance, egalitarian CUF, Borda):

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

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Examples (Full independance, egalitarian CUF, Borda):

m	$n = 2$	$n = 3$
4	ABBA	
5		
6		
8		
10		

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Examples (Full independance, egalitarian CUF, Borda):

m	$n = 2$	$n = 3$
4	ABBA	ABCC
5		
6		
8		
10		

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Examples (Full independance, egalitarian CUF, Borda):

m	$n = 2$	$n = 3$
4	ABBA	ABCC
5	AABBB	ABCCB
6		
8		
10		

Results

- **Full correlation**:
  - Utilitarian: trivial (every sequence is optimal)
  - Egalitarian: **NP**-complete (actually pseudo-polynomial), but very easy for QI
- **Full independence**:
  - Utilitarian + Borda: the alternating sequence (*ABABABAB*...) is optimal [Kalinowski et al., 2013]
  - Other cases: still open

Examples (Full independance, egalitarian CUF, Borda):

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

About manipulation...

Some (annoying?) feature...

About manipulation...

Some (annoying?) feature...

Example

2 agents (A, B), 4 objects:

- A: $o_1 \succ o_2 \succ o_3 \succ o_4$
- B: $o_2 \succ o_3 \succ o_4 \succ o_1$

Sequence $\sigma = ABBA \rightarrow \{o_1 o_4 | o_2 o_3\}$.

About manipulation...

Some (annoying?) feature...

Example

2 agents (A, B), 4 objects:

- A: $o_1 \succ o_2 \succ o_3 \succ o_4$
- B: $o_2 \succ o_3 \succ o_4 \succ o_1$

Sequence $\sigma = ABBA \rightarrow \{o_1 o_4 | o_2 o_3\}$.

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

About manipulation...

Some (annoying?) feature...

Example

2 agents (A, B), 4 objects:

- A: $o_1 \succ o_2 \succ o_3 \succ o_4$
- B: $o_2 \succ o_3 \succ o_4 \succ o_1$

Sequence $\sigma = ABBA \rightarrow \{o_1 o_4 | o_2 o_3\}$.

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

She can **manipulate** by picking o_2 instead of o_1 at first step $\rightarrow$ $\{o_1 o_2 | o_3 o_4\}$.

About manipulation...

Picking sequences are manipulable... How to prevent this?

Two approaches:

1. Computational barriers to manipulation...

About manipulation...

Picking sequences are manipulable... How to prevent this?

Two approaches:

1. Computational barriers to manipulation...
   - General manipulation: **polynomial** for 2 agents [B. and Lang, ECAI'14] but **NP-complete** for ≥ 3 agents [Aziz, B., Lang and Mackenzie, AAAI'17]

About manipulation...

Picking sequences are manipulable... How to prevent this?

Two approaches:

1. Computational barriers to manipulation...
   - General manipulation: **polynomial** for 2 agents [B. and Lang, ECAI'14] but **NP-complete** for ≥ 3 agents [Aziz, B., Lang and Mackenzie, AAAI'17]
   - Coalitional manipulation: **NP-complete** in general [B. and Lang, ECAI'14]

About manipulation...

Picking sequences are manipulable... How to prevent this?

Two approaches:

1. Computational barriers to manipulation...
   - General manipulation: **polynomial** for 2 agents [B. and Lang, ECAI'14] but **NP-complete** for ≥ 3 agents [Aziz, B., Lang and Mackenzie, AAAI'17]
   - Coalitional manipulation: **NP-complete** in general [B. and Lang, ECAI'14]
2. Strategyproof picking sequences...

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \dots AB \dots BC \dots C$).

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \ldots AB \ldots BC \ldots C$).

At first sight, **non-interleaving** sequences seem hardly compatible with fairness...

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \dots AB \dots BC \dots C$).

At first sight, **non-interleaving** sequences seem hardly compatible with fairness...

Is there a way to reconcile strategyproofness and fairness?

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \ldots AB \ldots BC \ldots C$).

At first sight, **non-interleaving** sequences seem hardly compatible with fairness...

Is there a way to reconcile strategyproofness and fairness?

For instance, for 3 agents, 10 objects, we "feel" that: $AABBBCCCCC$ is fairer than $AAAAABBBCC$

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \ldots AB \ldots BC \ldots C$).

At first sight, **non-interleaving** sequences seem hardly compatible with fairness...

Is there a way to reconcile strategyproofness and fairness?

For instance, for 3 agents, 10 objects, we "feel" that: $AABBBCCCCC$ is fairer than $AAAAABBBCC$

$\rightarrow$ We can compensate late arrival by higher number of goods picked.

Of strategyproof sequences

(Folk?) theorem

The only strategyproof picking sequences are those made of contiguous blocks of agents (e.g. $A \ldots AB \ldots BC \ldots C$).

At first sight, **non-interleaving** sequences seem hardly compatible with fairness...

Is there a way to reconcile strategyproofness and fairness?

For instance, for 3 agents, 10 objects, we "feel" that: *AABBBCCCCC* is fairer than *AAAAABBBCC*

→ We can compensate late arrival by higher number of goods picked.

Question

What is the *fairest* sequence?

Results

Good news [B., Gilbert, Lang and Méroué, arXiV'23]...

Proposition

For FI, FC, any $sw \in \{ut, eg, Na\}$ and any g, we can find an optimal sequence in time $O(m^2 \max(n, m))$ (dynamic programming)

Results

Good news [B., Gilbert, Lang and Méroué, arXiV'23]...

Proposition

For FI, FC, any $sw \in \{ut, eg, Na\}$ and any g, we can find an optimal sequence in time $O(m^2 \max(n, m))$ (dynamic programming)

Examples (Full independance, egalitarian CUF, Borda):

n	m	$sw = eg$	$sw = ut$
3	35	$(9, 10, 16)$	$(13, 11, 11)$
5	70	$(12, 12, 12, 13, 21)$	$(18, 16, 14, 11, 11)$
8	20	$(2, 2, 2, 2, 2, 3, 3, 4)$	$(3, 3, 3, 3, 2, 2, 2, 2)$
8	100	$(11, 11, 11, 11, 11, 12, 13, 20)$	$(18, 16, 15, 13, 12, 10, 8, 8)$

Results

Good news [B., Gilbert, Lang and Méroué, arXiV'23]...

Proposition

For FI, FC, any $sw \in \{ut, eg, Na\}$ and any g, we can find an optimal sequence in time $O(m^2 \max(n, m))$ (dynamic programming)

Discussion:

- Interest beyond picking sequences: under mild conditions, the only deterministic strategyproof mechanisms are within the family of **serial dictatorships** [Pápai, 2000, Pápai, 2001]
- Non-interleaving picking sequences $\approx$ a way to reconcile **strategyproofness**, (ex-ante) **fairness**, and (a form of) **efficiency**

Sequenceability as efficiency

Speaking of efficiency...

Sequenceability as efficiency

Speaking of efficiency...

- Usual way to characterize efficiency: **Pareto-efficiency**

Sequenceability as efficiency

Speaking of efficiency...

- Usual way to characterize efficiency: **Pareto-efficiency**
- Allocations obtained by picking sequences have a (weak) form of efficiency, that we call **sequenceability**

Sequenceability as efficiency

Speaking of efficiency...

- Usual way to characterize efficiency: **Pareto-efficiency**
- Allocations obtained by picking sequences have a (weak) form of efficiency, that we call **sequenceability**

Proposition [B. and Lemaître, COMSOC'16]

$\overrightarrow{\pi}$ is Pareto-efficient $\Rightarrow$ $\overrightarrow{\pi}$ is sequenceable

Incidentally, we also have:

- $\overrightarrow{\pi}$ is CEEI $\not\Rightarrow$ $\overrightarrow{\pi}$ is Pareto-efficient
- $\overrightarrow{\pi}$ is CEEI $\Rightarrow$ $\overrightarrow{\pi}$ is sequenceable

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation
- Let the agents **negotiate** (that is... trade bundles)

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation
- Let the agents **negotiate** (that is... trade bundles)

A particular kind of negotiation scheme: **trading cycles**

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation
- Let the agents **negotiate** (that is... trade bundles)

A particular kind of negotiation scheme: **trading cycles**

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation
- Let the agents **negotiate** (that is... trade bundles)

A particular kind of negotiation scheme: **trading cycles**

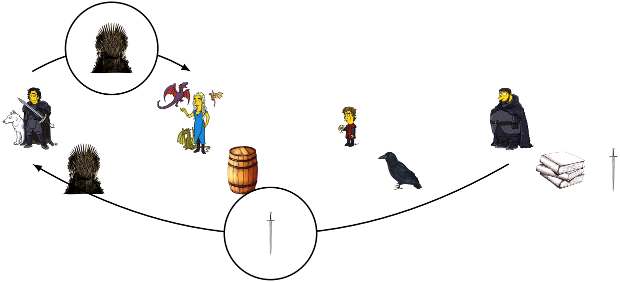

Swap deals vs sequences

Remember the third method to allocate indivisible goods? **Negotiation**...

- Start from an initial allocation
- Let the agents **negotiate** (that is... trade bundles)

A particular kind of negotiation scheme: **trading cycles**

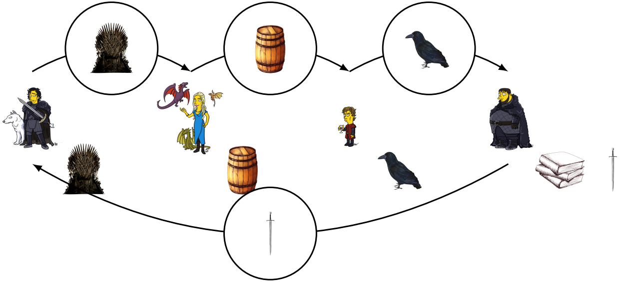

Trading cycles

(N, M)-cycle deal:

- N: cycle length
- M: max number of objects involved in each trade

(in the example before, $N = 4$ and $M = 1$)

Trading cycles

(N, M)-cycle deal:

- N: cycle length
- M: max number of objects involved in each trade

(in the example before, $N = 4$ and $M = 1$)

Proposition

$\overrightarrow{\pi}$ $(n, 1)$-cycle optimal $\Leftrightarrow$ $\overrightarrow{\pi}$ sequenceable. [Beynier, B., Lemaître, Maudet, Rey and Shams, AAMAS'19]

Trading cycles

(N, M)-cycle deal:

- N: cycle length
- M: max number of objects involved in each trade

(in the example before, $N = 4$ and $M = 1$)

Proposition

$\overrightarrow{\pi}$ $(n, 1)$-cycle optimal $\Leftrightarrow$ $\overrightarrow{\pi}$ sequenceable. [Beynier, B., Lemaître, Maudet, Rey and Shams, AAMAS'19]

Hence, (N, M)-cycle deals define:

- a **hierarchy of efficiency properties**
- whose highest level is **sequenceability**

Trading cycles

(N, M)-cycle deal:

- N: cycle length
- M: max number of objects involved in each trade

(in the example before, $N = 4$ and $M = 1$)

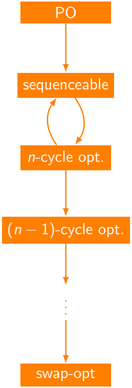

Proposition

$\overrightarrow{\pi}$ $(n, 1)$-cycle optimal $\Leftrightarrow$ $\overrightarrow{\pi}$ sequenceable. [Beynier, B., Lemaître, Maudet, Rey and Shams, AAMAS'19]

Hence, (N, M)-cycle deals define:

- a **hierarchy of efficiency properties**
- whose highest level is **sequenceability**

The full landscape of fairness

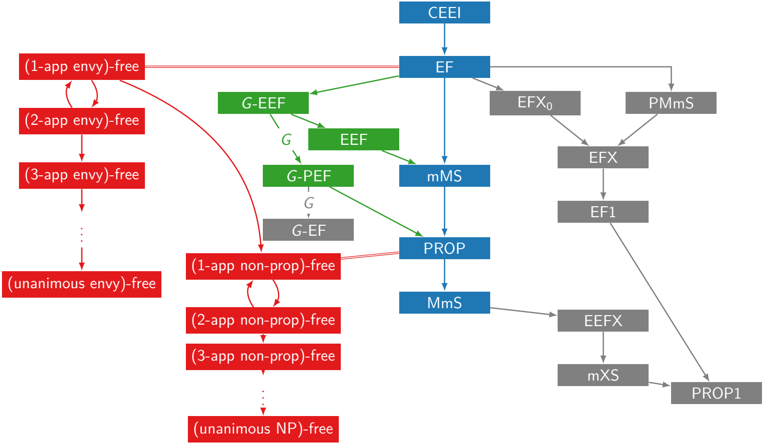

The full landscape of fairness

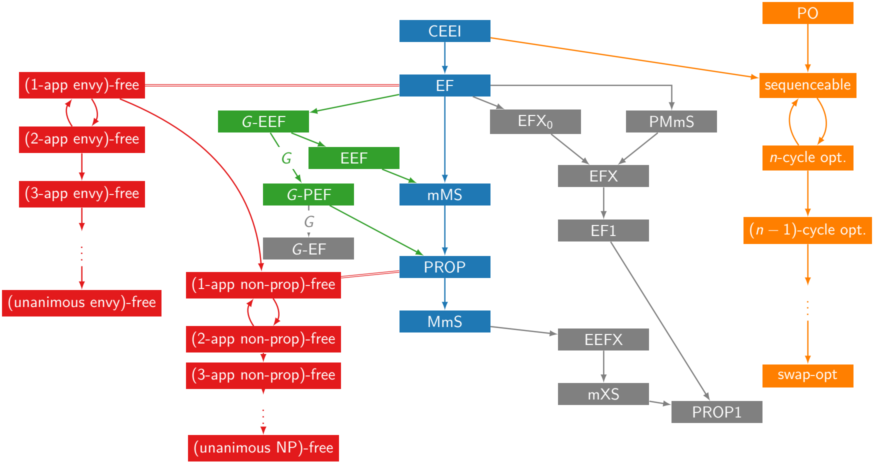

Voting

And the winner is... Alternative (fairer?) voting rules

From theory to experiments...

- So far, we have designed (supposedly) fair collective decision making procedures and studied their theoretical properties

From theory to experiments...

- So far, we have designed (supposedly) fair collective decision making procedures and studied their theoretical properties
- If we want to test how they behave **in practice**...

1. ...run lab experiments (with real humans)
2. ...run real-world experiments (with real humans as well)
3. ...run computer simulations (if possible, applied to real-world data)

From theory to experiments...

- So far, we have designed (supposedly) fair collective decision making procedures and studied their theoretical properties
- If we want to test how they behave **in practice**...

1. ...run lab experiments (with real humans)
2. ...run real-world experiments (with real humans as well)
3. ...run computer simulations (if possible, applied to real-world data)

From theory to experiments...

- So far, we have designed (supposedly) fair collective decision making procedures and studied their theoretical properties
- If we want to test how they behave **in practice**...
  1. ...run lab experiments (with real humans)
  2. ...run real-world experiments (with real humans as well)
  3. ...run computer simulations (if possible, applied to real-world data)
- In this part, we will focus on **voting** in the context of **political elections**
  - Experimenting alternative voting rules for the French presidential election
  - Simulating alternative voting rules for the French legislative election

From theory to experiments...

- So far, we have designed (supposedly) fair collective decision making procedures and studied their theoretical properties
- If we want to test how they behave **in practice**...
  1. ...run lab experiments (with real humans)
  2. ...run real-world experiments (with real humans as well)
  3. ...run computer simulations (if possible, applied to real-world data)
- In this part, we will focus on **voting** in the context of **political elections**
  - Experimenting alternative voting rules for the French presidential election
  - Simulating alternative voting rules for the French legislative election

Experimental setting

- An experiment run during the 2017 presidential election

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

- In France, the president is elected using **plurality with runoff**

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

- In France, the president is elected using **plurality with runoff**
- Many other rules exist (some of them known by SC theorists for centuries!)

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

- In France, the president is elected using **plurality with runoff**
- Many other rules exist (some of them known by SC theorists for centuries!)

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

- In France, the president is elected using **plurality with runoff**
- Many other rules exist (some of them known by SC theorists for centuries!)

Main question

How does the use of an alternative voting rule change the result of the election?

Experimental setting

- An experiment run during the 2017 presidential election
- Involving 10 researchers in France (economics + CS) and dozens of volunteers
- Online + *in situ* (5 places in France)
- 37739 participants online + 6358 *in situ* (incl. 1080 in Grenoble)

- In France, the president is elected using **plurality with runoff**
- Many other rules exist (some of them known by SC theorists for centuries!)

Main question

How does the use of an alternative voting rule change the result of the election?

Other similar experiments
[Baujard et al., 2014, Darmann et al., 2017, Darmann and Klamler, 2023]

More concretely...

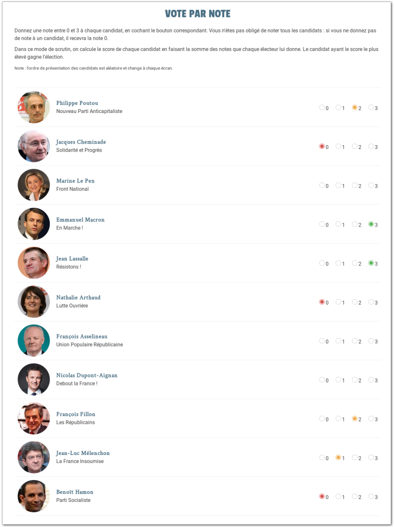

VOTE PAR NOTE

Donnez une note entre 0 et 3 à chaque candidat, en cochant le bouton correspondant. Vous n'êtes pas obligé de noter tous les candidats : si vous ne donnez pas de note à un candidat, il recevra la note 0.

Dans ce mode de scrutin, on calcule le score de chaque candidat en faisant la somme des notes que chaque électeur lui donne. Le candidat ayant le score le plus élevé gagne l'élection.

Note : l'ordre de présentation des candidats est aléatoire et change à chaque écran.

Philippe Poutou	Nouveau Parti Anticapitaliste	0	1	2	3
Jacques Cheminade	Solidarité et Progrès	0	1	2	3
Marine Le Pen	Front National	0	1	2	3
Emmanuel Macron	En Marche !	0	1	2	3
Jean Lassalle	Résistons !	0	1	2	3
Nathalie Arthaud	Lutte Ouvrière	0	1	2	3
François Asselineau	Union Populaire Républicaine	0	1	2	3
Nicolas Dupont-Aignan	Debout la France !	0	1	2	3
François Fillon	Les Républicains	0	1	2	3
Jean-Luc Mélenchon	La France Insoumise	0	1	2	3
Benoît Hamon	Parti Socialiste	0	1	2	3

More concretely...

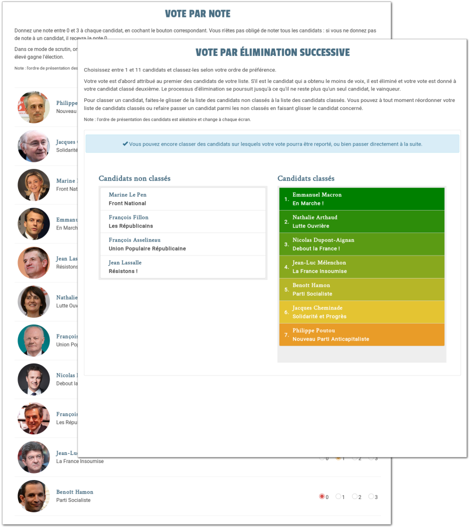

More concretely...

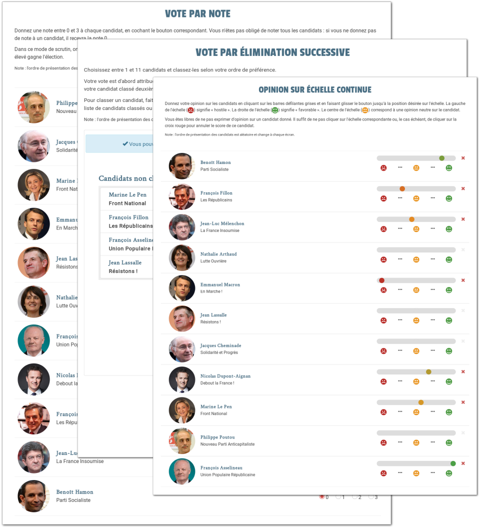

More concretely...

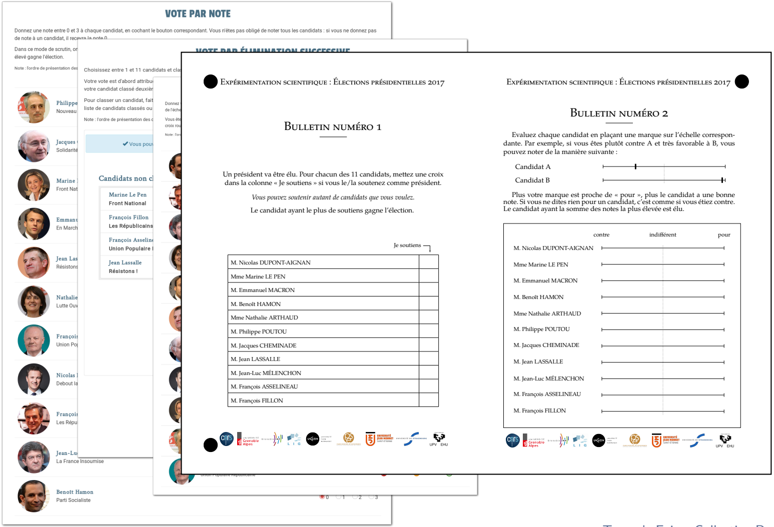

Results

Online experiment (corrected results)

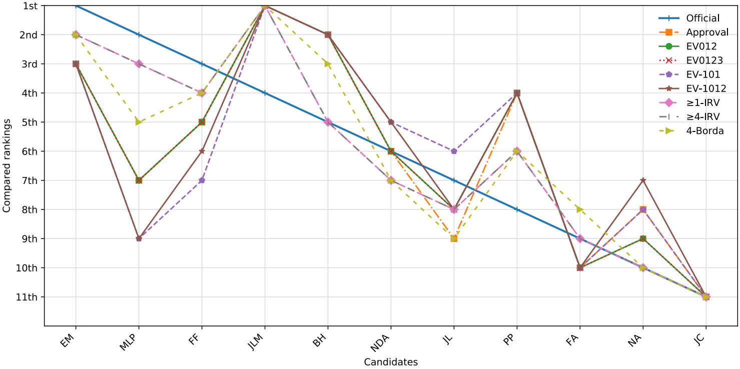

Results

In-situ experiment, in Grenoble (corrected results)

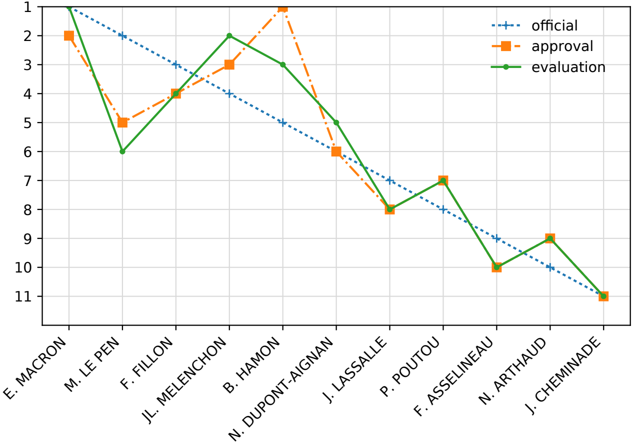

Results: discussion

- The results vary with the rules

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

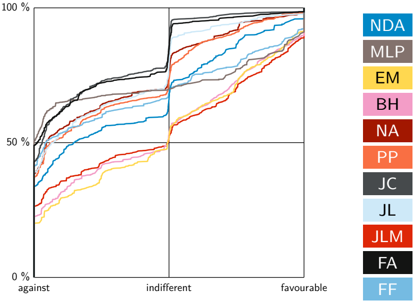

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

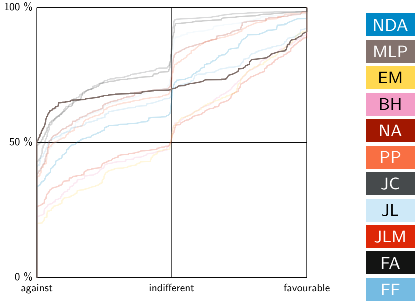

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

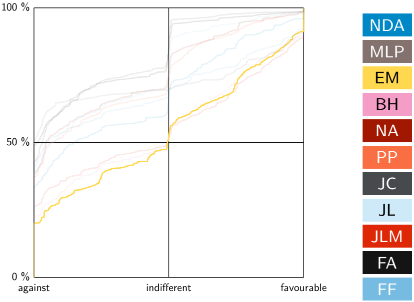

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

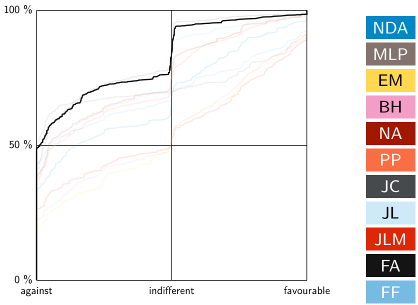

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

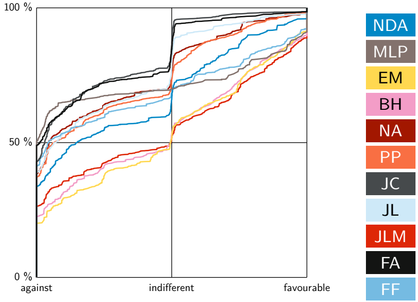

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

- Perception of the voting rules (alternative rules appreciated)

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

- Perception of the voting rules (alternative rules appreciated)
- Two datasets produced and published [B., Blanch, Baujard, Durand, Igersheim, Lang, Laruelle, Laslier, Lebon and Merlin, Zenodo'18 and 19]

Results: discussion

- The results vary with the rules
- Very biased population sample! → hard to unbias
- Several families of voting rules (official+IRV / Borda / AV+EV)
- Several kinds of candidates: polarizing, consensual, "small"

- Perception of the voting rules (alternative rules appreciated)
- Two datasets produced and published [B., Blanch, Baujard, Durand, Igersheim, Lang, Laruelle, Laslier, Lebon and Merlin, Zenodo'18 and 19]
- Part of the experiment run again in 2022

Conclusion and perspectives

A fair and safe operating space for humanity...

Conclusion

Fixme: TBC

Thank you

Want to know more?

http://recherche.noiraudes.net/en/hdr.php

Pictures borrowed from: https://drawthesimpsons.tumblr.com/

Akrami, H. and Garg, J. (2024).
Breaking the 3/4 barrier for approximate maximin share.
In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 74–91. SIAM.

Amanatidis, G., Markakis, E., Nikzad, A., and Saberi, A. (2017).
Approximation algorithms for computing maximin share allocations.
ACM Trans. Algorithms, 13(4).

Aziz, H., Bouveret, S., Caragiannis, I., Giagkousi, I., and Lang, J. (2018).
Knowledge, fairness, and social constraints.
In *Proceedings of the 32nd AAAI conference on Artificial Intelligence (AAAI'18)*, New Orleans, Louisiana, USA. AAAI Press.

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)*, San Francisco, California, USA. AAAI Press.

Baujard, A., Gavrel, F., Igersheim, H., Laslier, J.-F., and Lebon, I. (2014).
Who's favored by evaluative voting? an experiment conducted during the 2012 french presidential election.
Electoral Studies, 34:131–145.

Beynier, A., Bouveret, S., Lemaître, M., Maudet, N., Rey, S., and Shams, P. (2019).
Efficiency, sequenceability and deal-optimality in fair division of indivisible goods.
In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'19)*, page 9 pages, Montreal, Canada. IFAAMAS.

Bouveret, S., Blanch, R., Baujard, A., Durand, F., Igersheim, H., Lang, J., Laruelle, A., Laslier, J.-F., Lebon, I., and Merlin, V. (2018).
Voter autrement 2017 - online experiment.
Dataset and companion article on Zenodo.

Bouveret, S., Blanch, R., Baujard, A., Durand, F., Igersheim, H., Lang, J., Laruelle, A., Laslier, J.-F., Lebon, I., and Merlin, V. (2019).
Voter Autrement 2017 for the French Presidential Election – The data of the In Situ Experiment.
Dataset and companion article on Zenodo.

Bouveret, S., Gilbert, H., Lang, J., and Méroué, G. (2023).
Thou shalt not pick all items if thou are first: of strategyproof and fair picking sequences.

Bouveret, S. and Lang, J. (2011).
A general elicitation-free protocol for allocating indivisible goods.
In *Proceedings of the 22st International Joint Conference on Artificial Intelligence (IJCAI'11)*, Barcelona, Spain.

Bouveret, S. and Lang, J. (2014).
Manipulating picking sequences.
In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, Prague, Czech Republic. IOS Press.

Bouveret, S. and Lemaître, M. (2014).
Characterizing conflicts in fair division of indivisible goods using a scale of criteria.
In *Proceedings of the 13th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS'14)*, Paris, France. ACM.

Bouveret, S. and Lemaître, M. (2015).
Characterizing conflicts in fair division of indivisible goods using a scale of criteria.
Autonomous Agents and Multi-Agent Systems, 30(2):259–290.

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)*, Toulouse, France.

Brânzei, S., Hosseini, H., and Miltersen, P. B. (2015).

Characterization and computation of equilibria for indivisible goods.
In Hoefer, M., editor, *Algorithmic Game Theory*, pages 244–255, Berlin, Heidelberg. Springer Berlin Heidelberg.

Budish, E. (2011).

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

Caragiannis, I., Garg, J., Rathi, N., Sharma, E., Varricchio, G., et al. (2023).

New fairness concepts for allocating indivisible items.
In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, pages 2554–2562.

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 2016 ACM Conference on Economics and Computation*, EC '16, pages 305–322, New York, NY, USA. ACM.

Darmann, A., Grundner, J., and Klamler, C. (2017).

Election outcomes under different ways to announce preferences: an analysis of the 2015 parliament election in the austrian federal state of styria.
Public Choice, 173:201–216.

Darmann, A. and Klamler, C. (2023).

Does the rule matter? a comparison of preference elicitation methods and voting rules based on data from an austrian regional parliamentary election in 2019.
Public Choice, 197(63–87).

Foley, D. K. (1967).

Resource allocation and the public sector.
Yale Econ. Essays, 7(1).

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

A social welfare optimal sequential allocation procedure.
In Rossi, F., editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, pages 227–233, Beijing, China. IJCAI/AAAI.

Kurokawa, D., Procaccia, A. D., and Wang, J. (2016).

When can the maximin share guarantee be guaranteed?
In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, page 523–529. AAAI Press.

Lipton, R. J., Markakis, E., Mossel, E., and Saberi, A. (2004).

On approximately fair allocations of indivisible goods.
In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131.

Parijs, P. V. (1997).

Real Freedom for All: What (if Anything) Can Justify Capitalism?
Number 9780198293576 in OUP Catalogue. Oxford University Press.

Procaccia, A. D. and Wang, J. (2014).

Fair enough: Guaranteeing approximate maximin shares.
In *Proceedings of the 15th ACM Conference on Economics and Computation (EC'14)*.

Pápai, S. (2000).

Original papers: Strategyproof multiple assignment using quotas.
Review of Economic Design, 5(1):91–105.

Pápai, S. (2001).

Strategyproof and nonbossy multiple assignments.
Journal of Public Economic Theory, 3(3):257–271.

Sandholm, T. W. (1998).

Contract types for satisficing task allocation: I. theoretical results.
In Sen, S., editor, *Proceedings of the AAAI Spring Symposium: Satisficing Models*, pages 68–75, Menlo Park, California. AAAI Press.

Shams, P., Beynier, A., Bouveret, S., and Maudet, N. (2021).
Minimizing and balancing envy among agents using ordered weighted average.
In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT'21)*, Toulouse, France. Springer Verlag.

Shams, P., Beynier, A., Bouveret, S., and Maudet, N. (2022).
Fair in the eyes of others.
Journal of Artificial Intelligence Research (JAIR), 75:913–951.

Shapley, L. and Scarf, H. (1974).
On cores and indivisibility.
Journal of mathematical economics, 1(1):23–37.

Steinhaus, H. (1948).
The problem of fair division.
Econometrica, 16(1):101–104.