Mathematical programming in computational social choice

Gerdus Benade

benade@bu.edu

Human factors
Strategic agents
Partial information
Biases
Fairness

Fair division

Social choice

Optimization

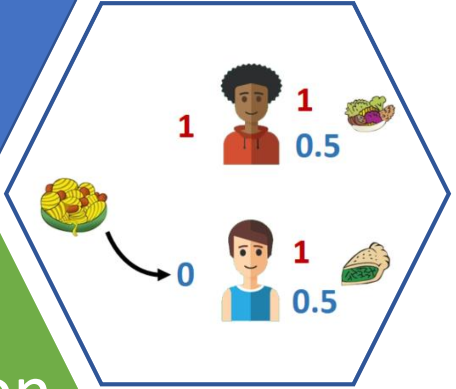

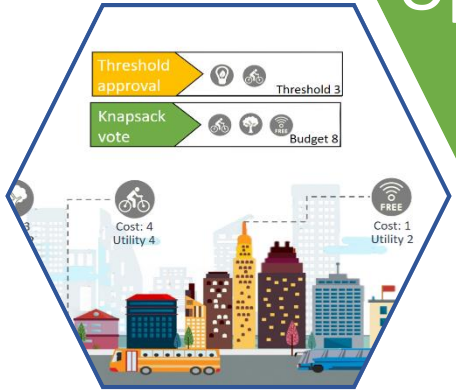

Preference elicitation for participatory budgeting

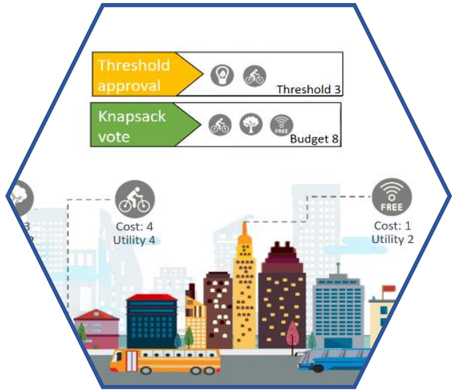

Gerdus Benade

Preference elicitation for participatory budgeting




Swaprava Nath   Ariel Procaccia   Nisarg Shah   Kobi Gal

Let's warm up with a budgeting problem

- Coffee, transport and relaxation
- Each provides a different value per dollar spent (v_c, v_r, v_t)
- Allocate x_c, x_r, x_t to each category
- Maximize utility gained from budget B

$$\max v_c x_c + v_r x_r + v_t x_t$$

$$\begin{aligned} s.t. \quad & x_r + x_c + x_t \leq B \\ & x_r, x_c, x_t \geq 0 \end{aligned}$$

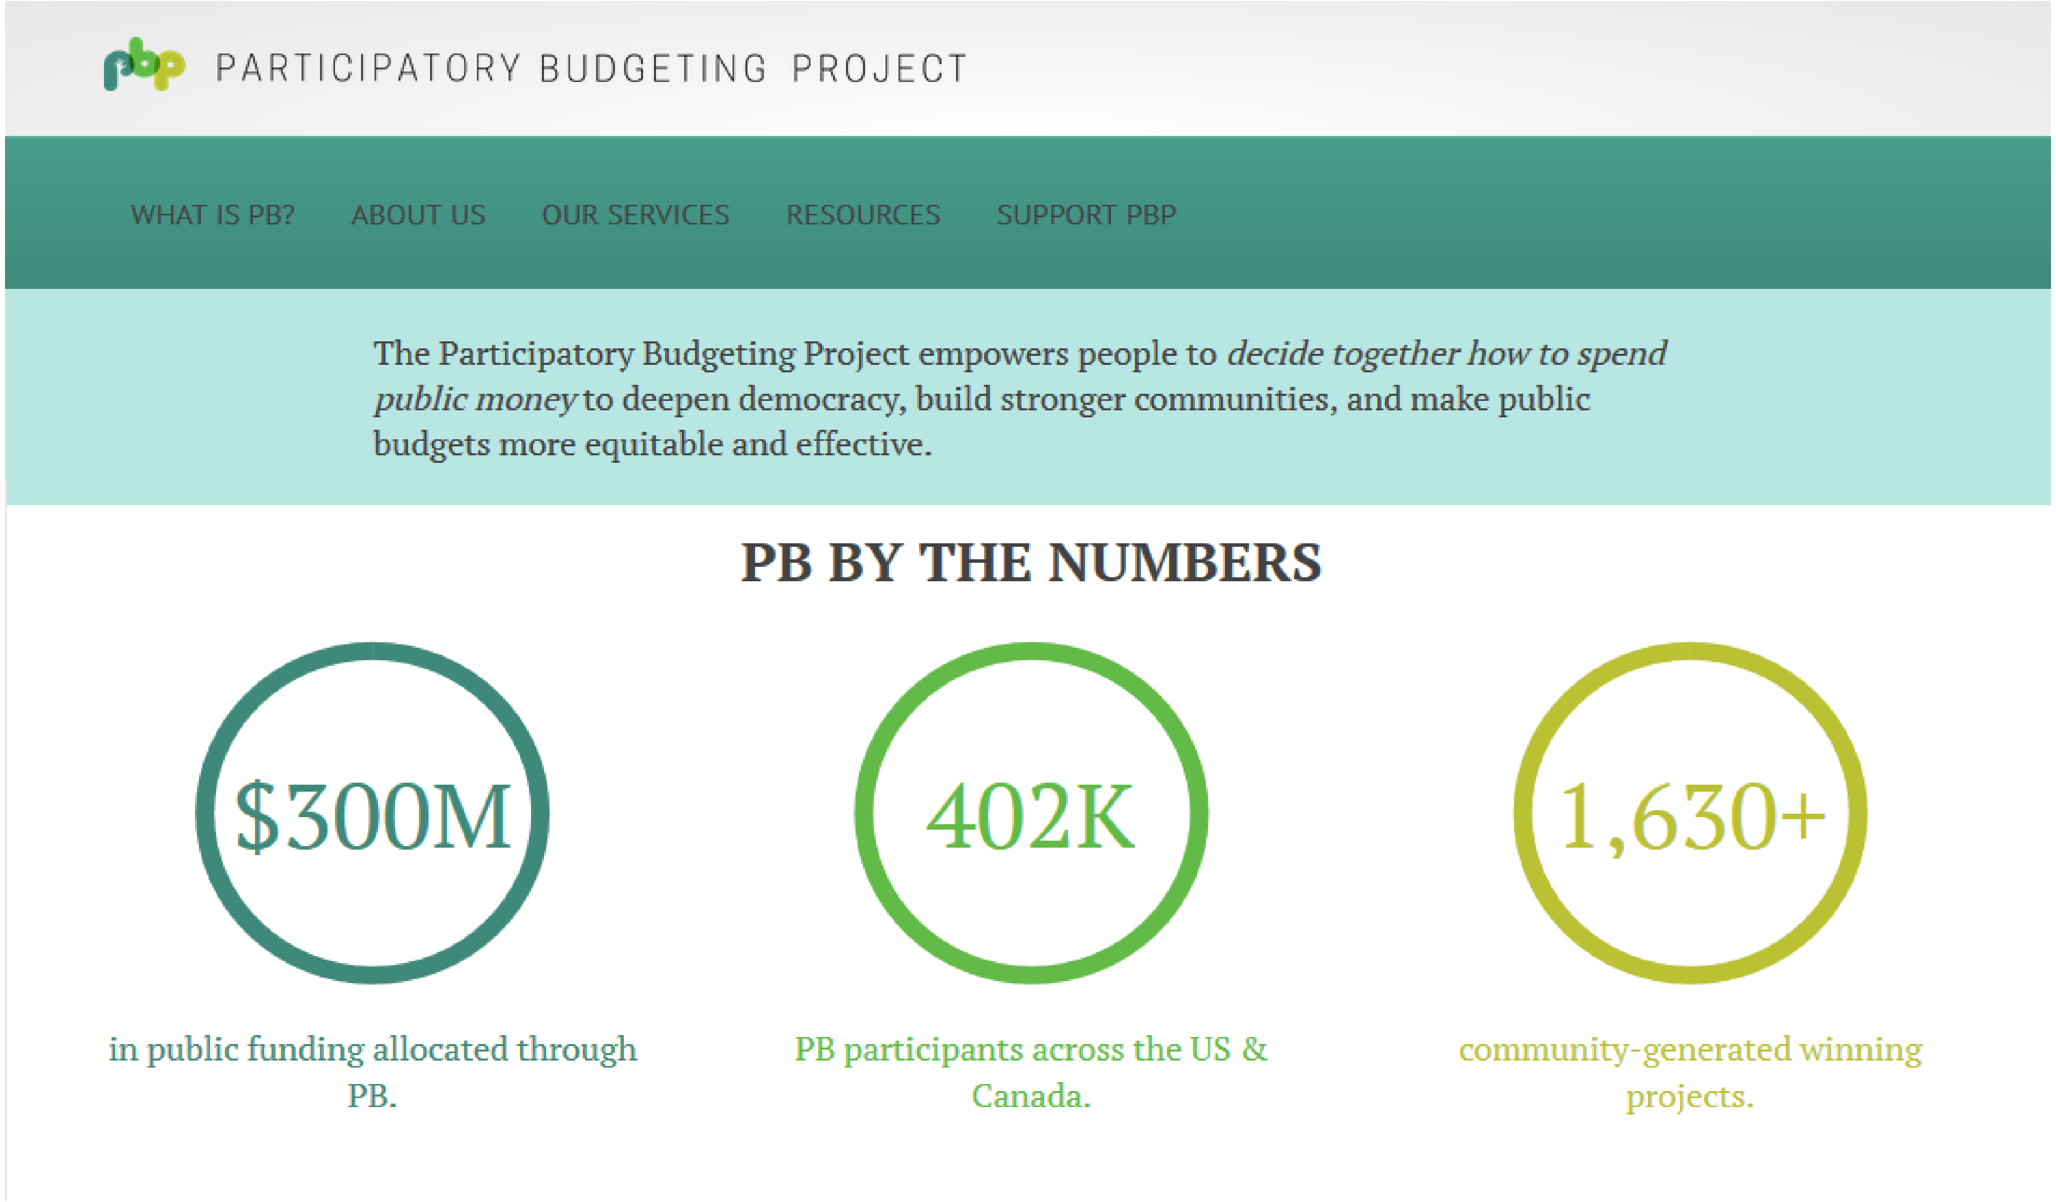

Image from http://participatorybudgeting.org https://pb.cambridgema.gov/

CITY OF CAMBRIDGE PARTICIPATORY BUDGETING

Home | PB 6 Final 20 | Events | Blog | FAQs | Get Involved | Project Status | PB Cycles | Contact

How should voters express their preferences?

How should the votes be aggregated?

~~Can we limit the cognitive burden on voters?~~

Model

- n residents N, m projects A
- Project a has cost c_a, global budget $B = 1$
- Voter i has utility function v_i
- Utilities are additive

$$v_i(S) = \sum_{a \in S} v_i(a) \text{ for every } S \subseteq A$$

Model

- n residents N, m projects A
- Project a has cost c_a, global budget $B = 1$
- Voter i has utility function v_i
- Utilities are additive and normalized

$$v_i(A) = 1 \text{ for all } i \in N$$

Model

- n residents N, m projects A
- Project a has cost c_a, global budget $B = 1$
- Voter i has utility function v_i, utility v_{ia} for alternative a
- Utilities are additive and normalized

$$\max_S \quad \mathrm{sw}(S) = \textstyle\sum_{a \in S} \sum_i v_i(a)$$

$$\text{subject to} \quad c(S) \leq B$$

Interlude: the knapsack problem

$$\left(\begin{array}{ll} \max\limits_{x} & \sum_a \left(\sum_i v_{ia}\right) x_a \\ \text{subject to} & \sum_a c_a x_a \le B \\ & \qquad x_a \in \{0,1\} \;\; \forall a \end{array}\right) \quad \text{KNAP}(v)$$

Interlude: the knapsack problem

$$\left(\begin{array}{ll} \max_x & \sum_a (\sum_i v_{ia}) x_a \\ \text{subject to} & \sum_a c_a x_a \leq B \\ & \quad x_a \in \{0,1\} \ \forall a \end{array}\right) \quad \text{KNAP}(v)$$

- Notice integer variables x_a
- LP solvable in polynomial time, IP takes exponential time in general
- Knapsack problems are ‘easy’: thousands of items takes seconds

Model

- n residents N, m projects A
- Project a has cost c_a, global budget $B = 1$
- Voter i has utility function v_i, utility v_{ia} for alternative a
- Utilities are additive and normalized

$$\max_S \quad \mathrm{sw}(S) = \sum_{a \in S} \sum_i v_i(a)$$

$$\text{subject to} \quad c(S) \leq B$$

- Assume voter i submits vote ρ_i consistent with v_i : $v_i \rhd \rho_i$

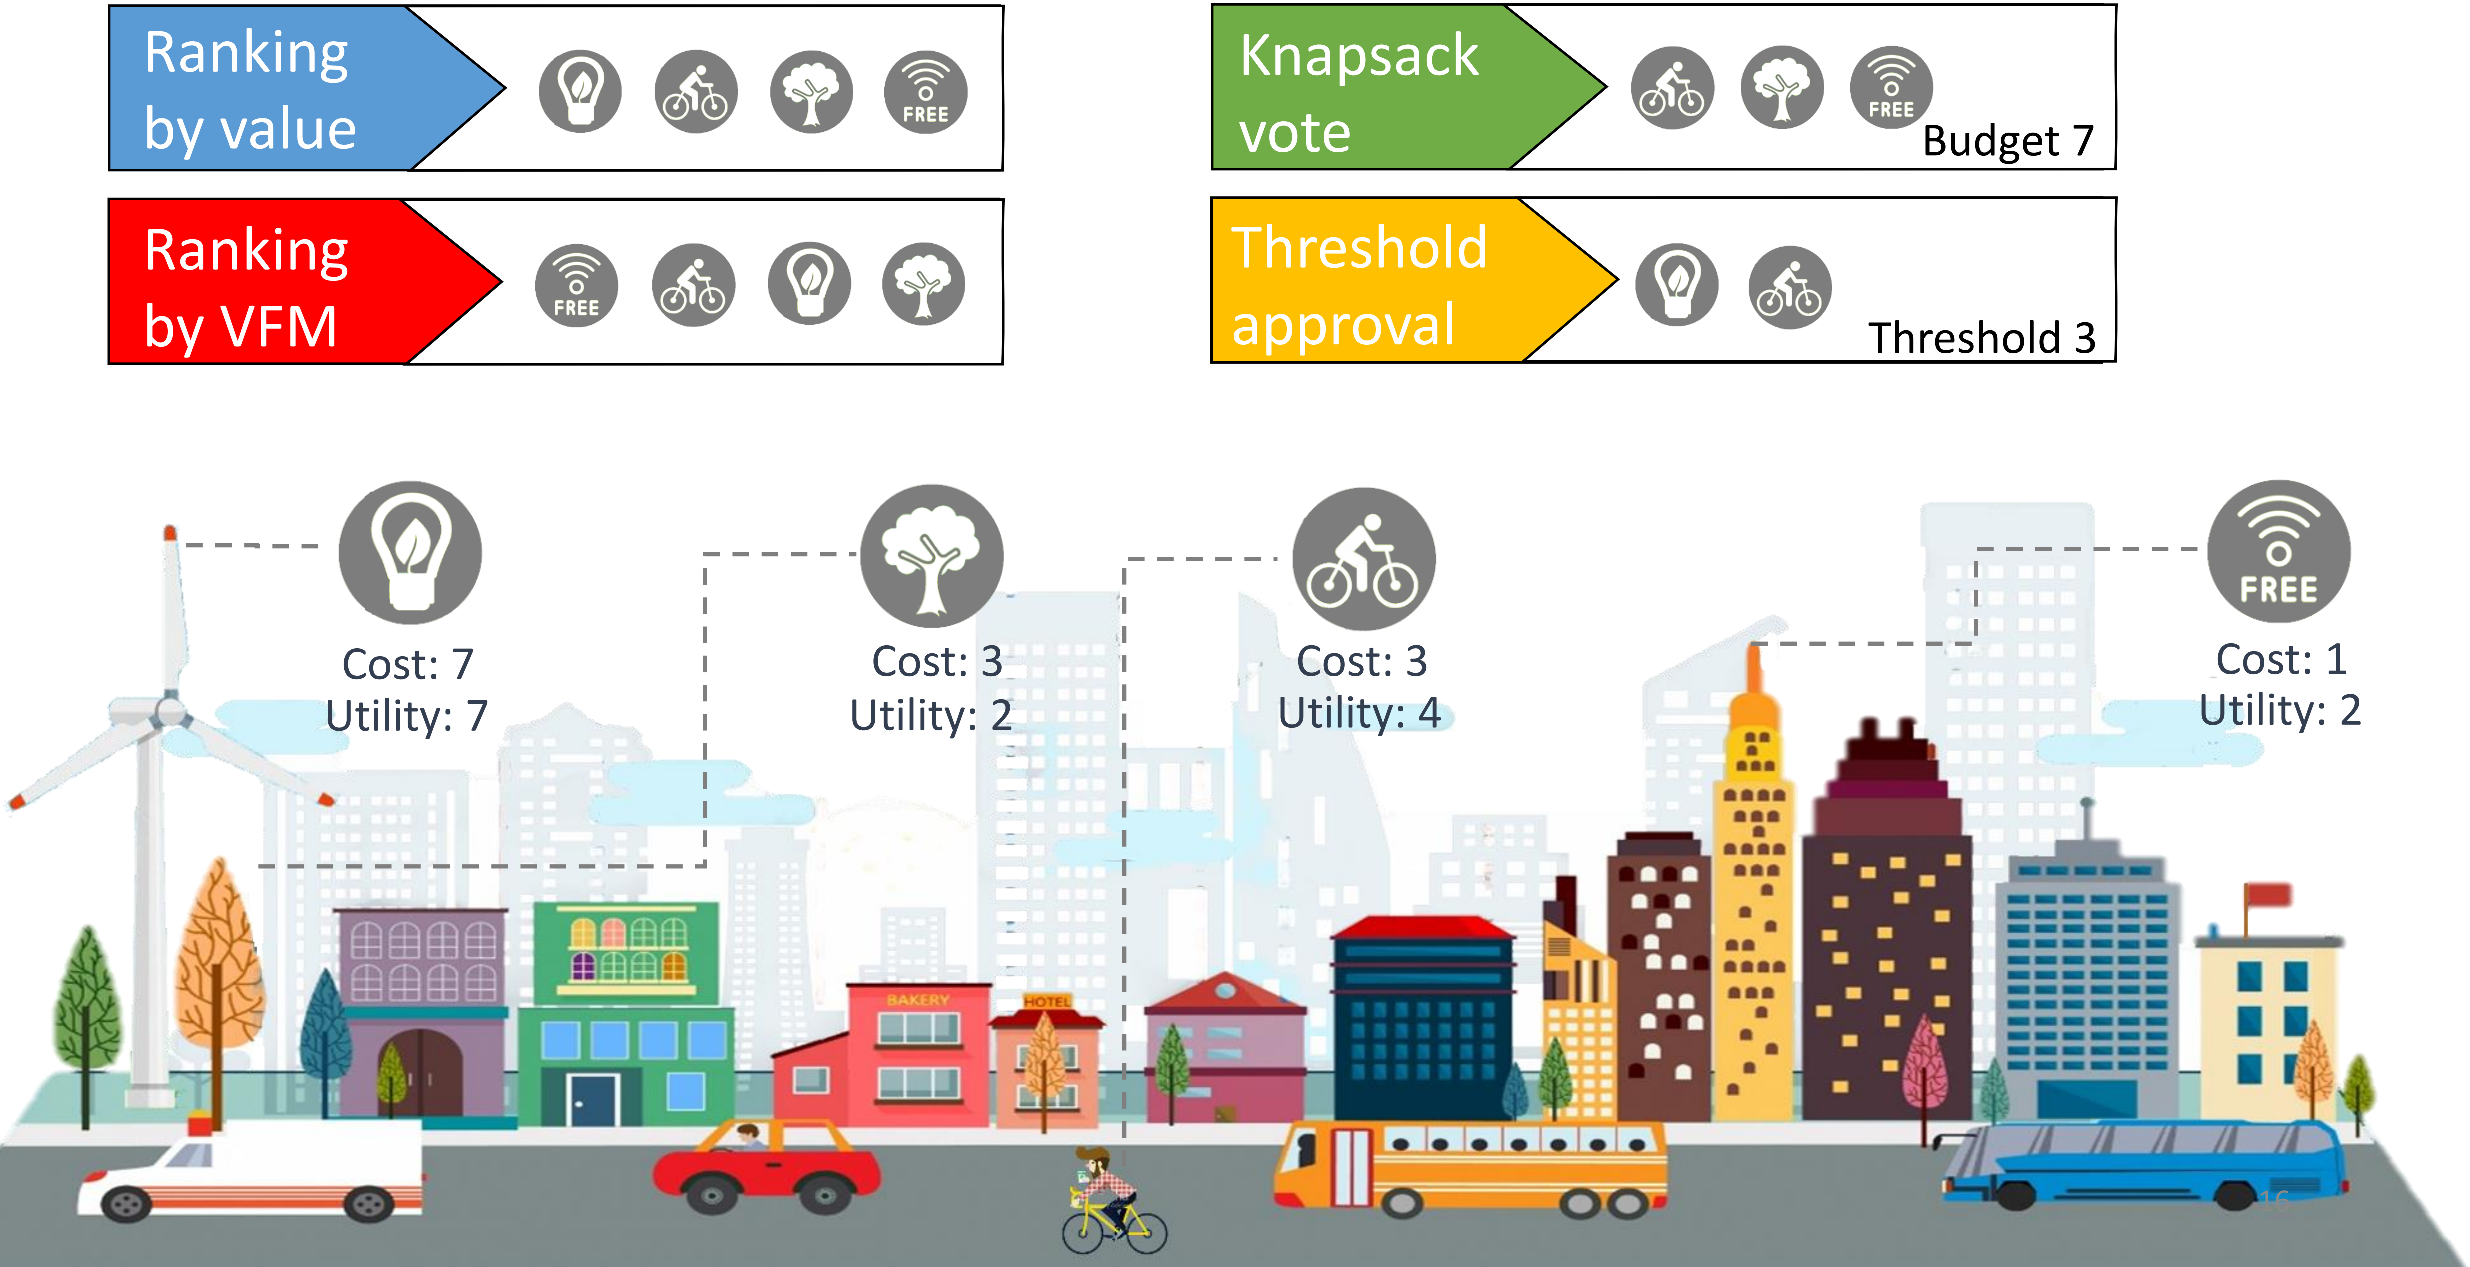
Ranking by value
Ranking by VFM
Knapsack vote
Budget 7
Threshold approval
Threshold 3
FREE
Cost: 7
Utility: 7
Cost: 3
Utility: 2
Cost: 3
Utility: 4
Cost: 1
Utility: 2
BAKERY
HOTEL

Implicit utilitarian voting

- Randomized voting rule f vote profile $\vec{\rho}$ as input, returns distribution over feasible sets of projects $f(\vec{\rho})$

Knapsack solution for $\vec{v}$

$$\mathrm{dist}(f,\vec{\rho}) = \underbrace{\max_{\vec{v}\,:\,\vec{v}\,\rhd\,\vec{\rho}}}_{\substack{\text{Worst case over}\\ \text{consistent utilities}}} \underbrace{\frac{\max_{T:c(T)\le B}\ sw(T,\vec{v})}{\mathbb{E}\left[sw(f(\vec{\rho}),\vec{v})\right]}}_{\substack{\text{Social welfare ratio}\\ \text{(Approximation ratio on } \vec{v}\text{)}}}$$

Social welfare of f on $\vec{v}$

Implicit utilitarian voting

- Randomized voting rule f vote profile $\vec{\rho}$ as input, returns distribution over feasible sets of projects $f(\vec{\rho})$

$$\mathrm{dist}(f, \vec{\rho}) = \max_{\vec{v} \,:\, \vec{v} \rhd \vec{\rho}} \frac{\max_{T: c(T) \leq B} \; sw(T, \vec{v})}{\mathbb{E}\left[sw(f(\vec{\rho}), \vec{v})\right]}$$

$$\mathrm{dist}(f) = \underbrace{\max_{\vec{\rho}}}_{\text{Worst case over inputs}} \max_{\vec{v} \,:\, \vec{v} \rhd \vec{\rho}} \frac{\max_{T: c(T) \leq B} \; sw(T, \vec{v})}{\mathbb{E}\left[sw(f(\vec{\rho}), \vec{v})\right]}$$

Decouple input format and aggregation

- Distortion of a voting rule:

$$\text{dist}(f) = \max_{\vec{\rho}} \max_{\vec{v} \,:\, \vec{v} \rhd \vec{\rho}} \frac{\max_{T:c(T)\leq B} \; sw(T, \vec{v})}{\mathbb{E}\left[sw(f(\vec{\rho}), \vec{v})\right]}$$

- Distortion of an input format: the distortion of its best voting rule

- Knapsack voting $\widetilde{\Theta}(m)$
- Ranking by value $\widetilde{\Theta}(\sqrt{m})$
- Ranking by VFM $\widetilde{\Theta}(\sqrt{m})$
- Threshold approval $O(\log^2 m)$

How should voters express their preferences?

How should the votes be aggregated?

Can we limit the cognitive burden on voters?

How should the votes be aggregated?

Current practice: greedy aggregation based on number of approvals/appearances in a knapsack etc. (Goel et al.)

Use the input-specific distortion-minimizing voting rules:

Deterministic: $$f^*(\vec{\rho}) = \underset{S}{\operatorname{argmin}} \max_{\vec{v}\,:\,\vec{v} \rhd \vec{\rho}} \frac{\max_{T:c(T)\leq B} \; sw(T,\vec{v})}{sw(S,\vec{v})}$$

Distortion-minimizing voting rules via LP

Deterministic: $f^*(\vec{\rho}) = \operatorname{argmin} \max_{\vec{v}\,:\,\vec{v} \rhd \vec{\rho}} \dfrac{\max_{T:c(T)\leq B}\ sw(T,\vec{v})}{sw(S,\vec{v})}$

For every S:

For every $\vec{v}\ :\ \vec{v}\ \rhd\ \vec{\rho}$:

Compute optimal solution to KNAP$(\vec{v}), T$

Keep track of worst ratio $\mathrm{sw}(\mathrm{T},\vec{v})/sw(S,\vec{v})$

Return S with smallest worst-case ratio

Distortion-minimizing voting rules via LP

Deterministic: $f^*(\vec{\rho}) = \operatorname{argmin} \max_{\vec{v} \,:\, \vec{v} \rhd \vec{\rho}} \dfrac{\max_{T:c(T)\leq B} \; sw(T,\vec{v})}{sw(S,\vec{v})}$

For every S:

For every T:

Compute $\max_{\vec{v} \,:\, \vec{v} \rhd \vec{\rho}} \mathrm{sw}(\mathrm{T}, \vec{v})/sw(S, \vec{v})$

Return S with smallest worst-case ratio

Inner LP (S, T fixed)

$$\max_{\vec{v}\,:\,\vec{v} \rhd \vec{\rho}} \mathrm{sw}(\mathrm{T}, \vec{v}) / sw(S, \vec{v})$$

- x_S, x_T are characteristic vectors of sets S, T (fixed)
  - $|A|$-dimensional
  - $x_S(a) = 1$ if alternative a is in set S
- Variables $v_i(a)$ for every voter i, alternative a

$$\max \frac{\sum_a \sum_i v_i(a) x_T(a)}{\sum_a \sum_i v_i(a) x_S(a)}$$

$$\text{s.t. } \vec{v} \geq 0\,, \vec{v} \rhd \vec{\rho}$$

Consistency constraints are linear

- For example, for a ranking ρ_i
  - $v_i\big(\rho_i(1)\big) \geq v_i\big(\rho_i(2)\big) \geq \cdots \geq v\_i(\rho_i(m))$
  - Non-negativity + normalization constraints

- For threshold approval votes with threshold t
  - $v_i(a) \geq t$ if a approved
  - $v_i(a) \leq t$ if not approved
  - Non-negativity + normalization constraints

Distortion-minimizing voting rules via LP

For every S:

For every T:

Solve $\max \dfrac{\sum_a \sum_i v_i(a) x_T(a)}{\sum_a \sum_i v_i(a) x_S(a)}$

s.t. $\vec{v} \geq 0 \,, \vec{v} \;\rhd\; \vec{\rho}$

LP with fractional objective:
Charnes-Cooper transformation

Return S with smallest worst-case ratio

Experiments

- Real-world participatory budgeting elections
  - Held in Cambridge (MA) in 2015 and 2016
  - Data provided by Ashish Goel and the Stanford Crowdsourced Democracy Team
- 10 projects, 4000 voters
- Real votes $\vec{\rho}$ → consistent utility $\vec{v}$ → votes in all formats → aggregate
- Measure social welfare ratio, compare 4 formats + greedy baselines

$$\left(\frac{\max sw(T, \vec{v})}{sw(f(\vec{\rho}), \vec{v})}\right)$$

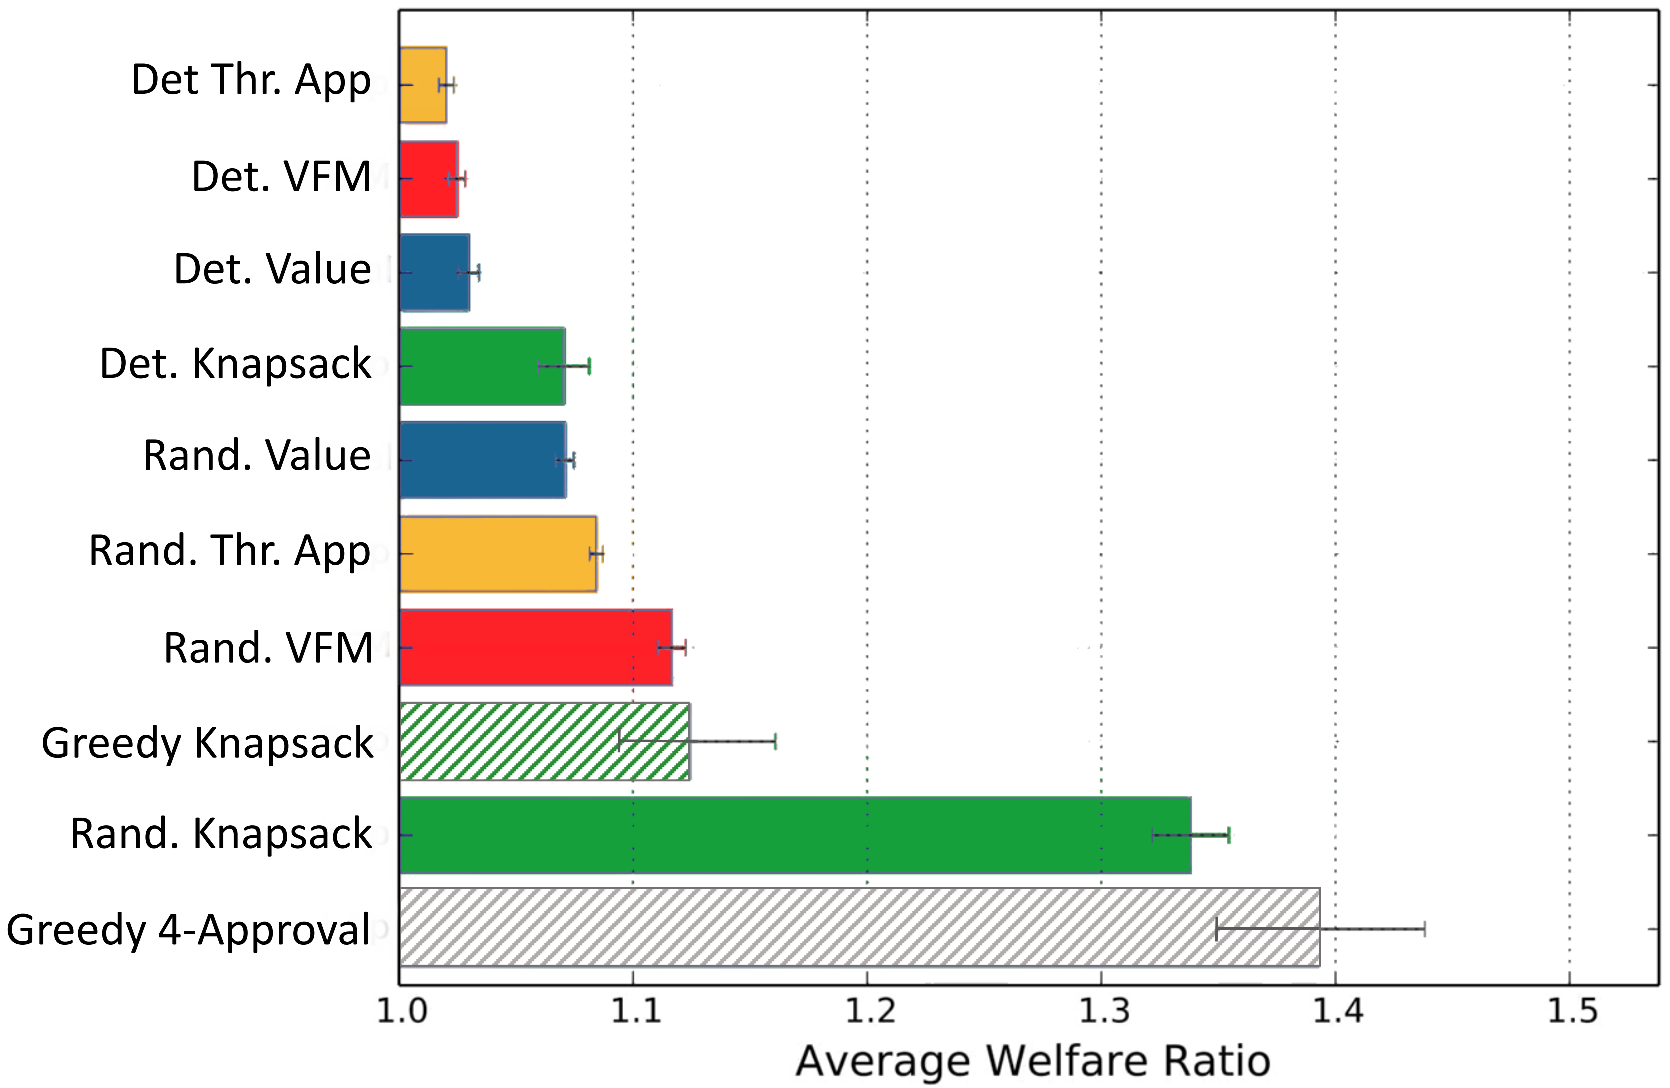
Det Thr. App
Det. VFM
Det. Value
Det. Knapsack
Rand. Value
Rand. Thr. App
Rand. VFM
Greedy Knapsack
Rand. Knapsack
Greedy 4-Approval
1.0
1.1
1.2
1.3
1.4
1.5
Average Welfare Ratio

Human factors
Strategic agents
Partial information
Biases
Fairness

Fair division

Social choice

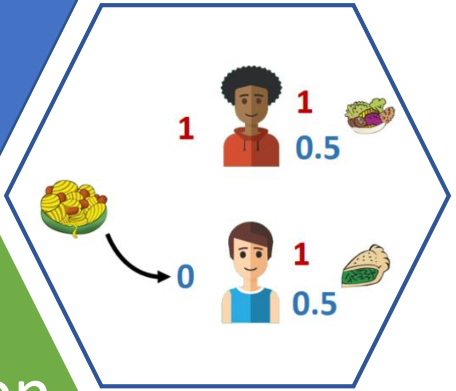

Optimization

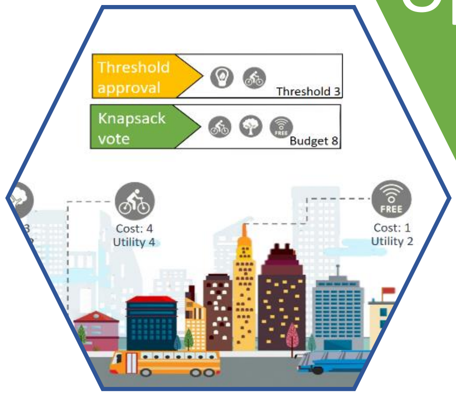

How to Make Envy Vanish over Time

Alex Psomas Aleks Kazachkov Ariel Procaccia

How can we fairly allocate goods that arrive over time?

How can we fairly allocate goods that arrive over time?

Envy-freeness

No agent prefers another allocation to their own.

$$ENVY^{ij} = \max\{v_i(A_j) - v_i(A_i), 0\}$$

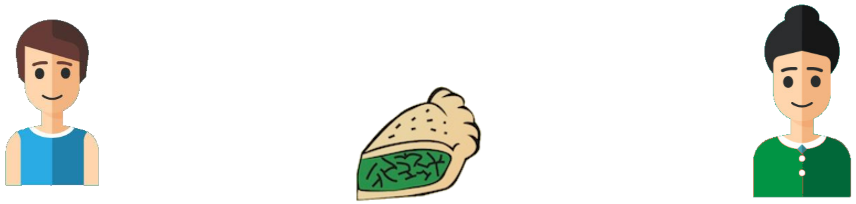

How can we fairly allocate goods that arrive over time?

Envy-freeness

No agent prefers another allocation to their own.

$$ENVY^{ij} = \max\{v_i(A_j) - v_j(A_i), 0\}$$

Indivisible goods: EF1 (Envy-free up to 1 good)

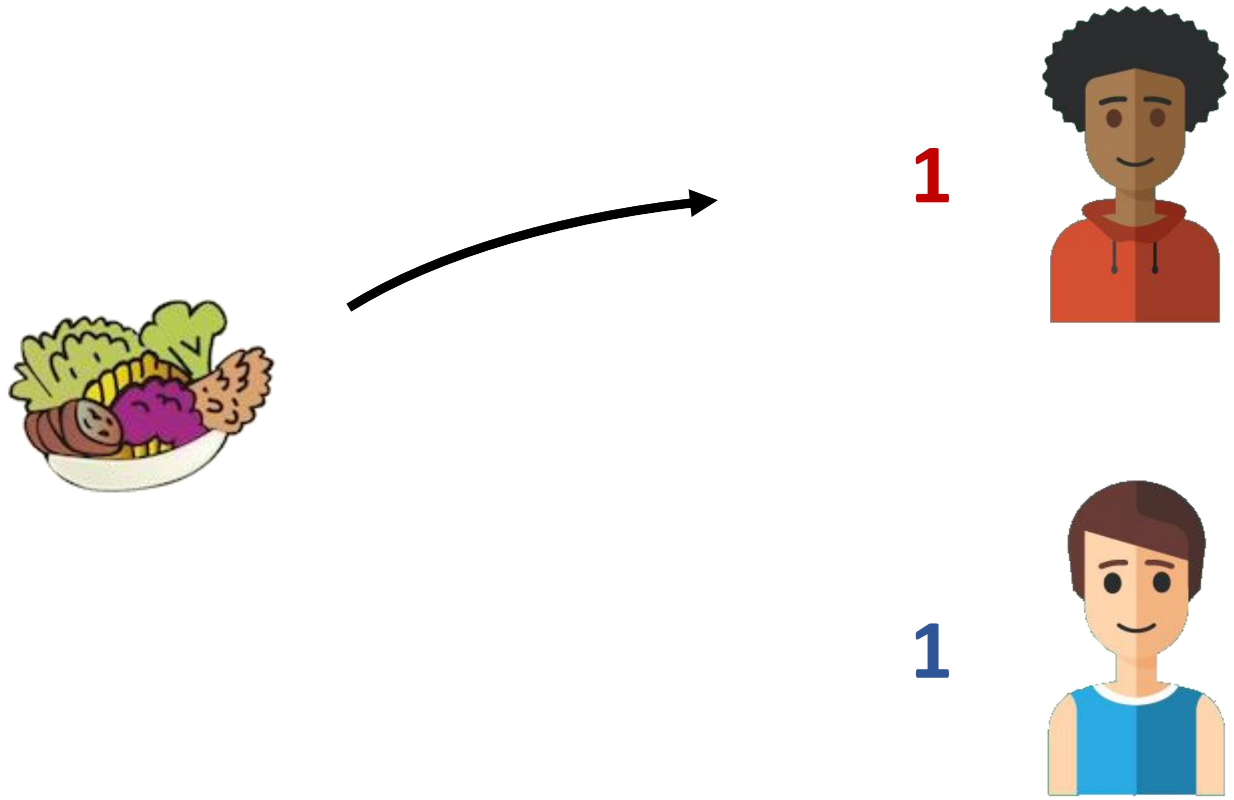
1
1

1

0

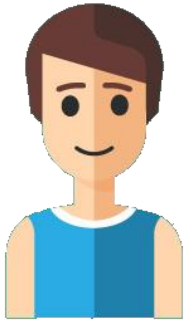

0

1

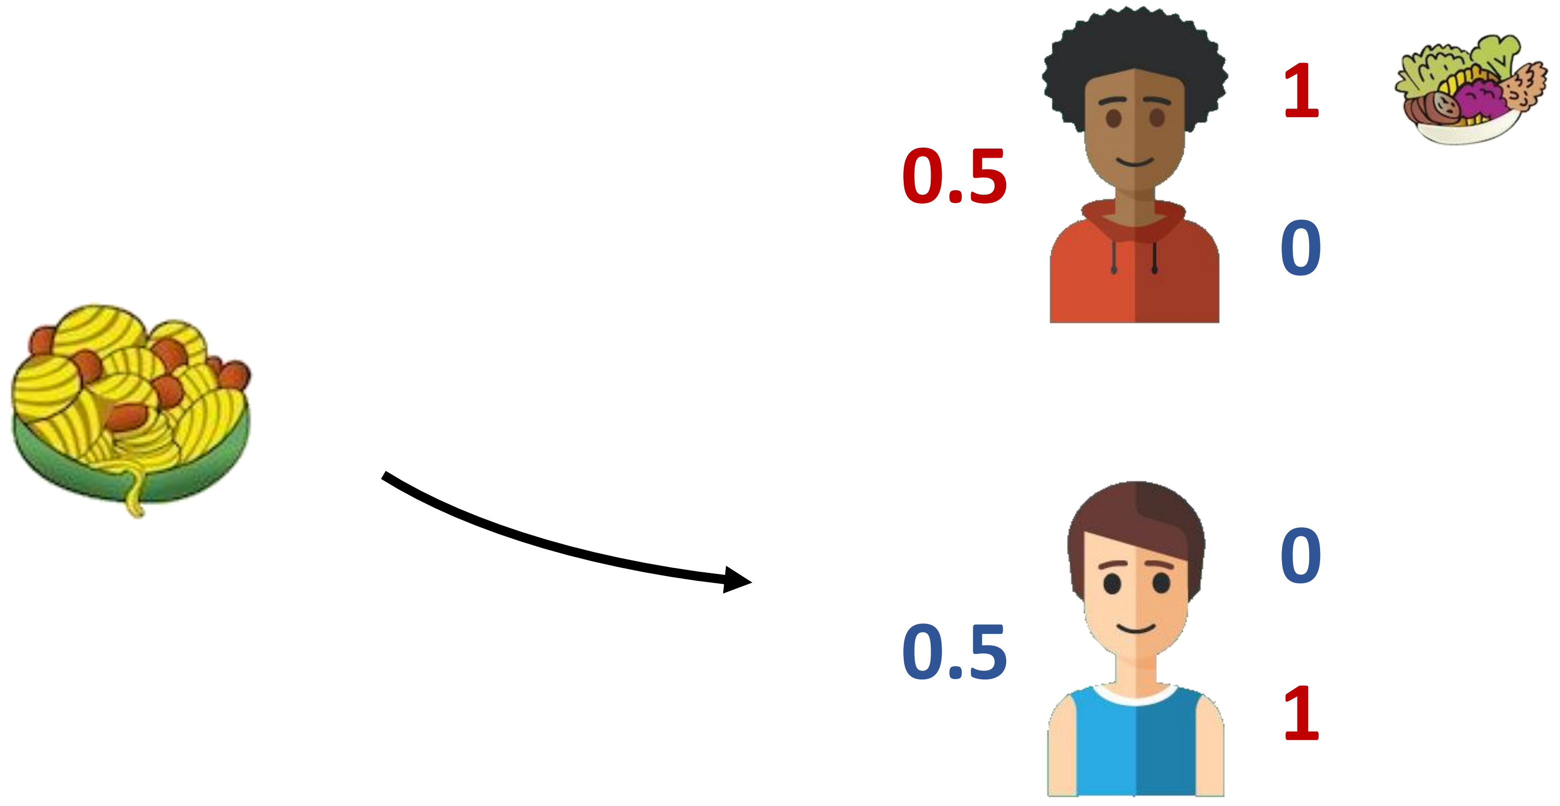
1
0.5
0
0
0.5
1

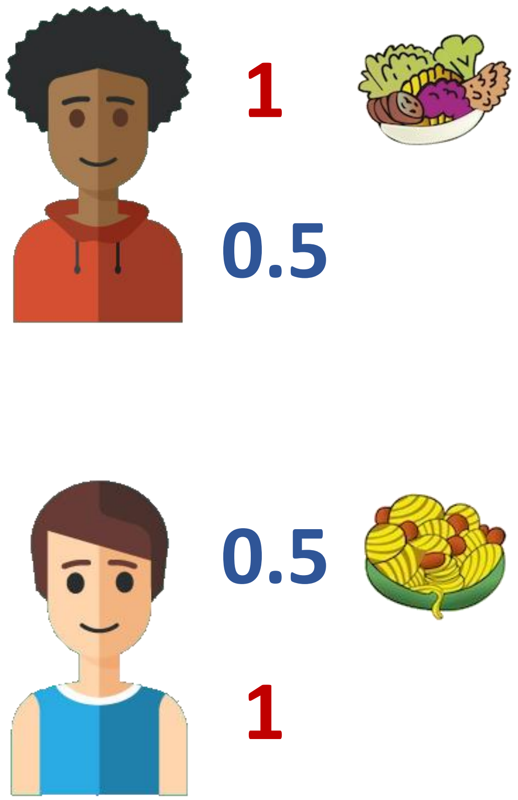
1
0.5
0.5
1

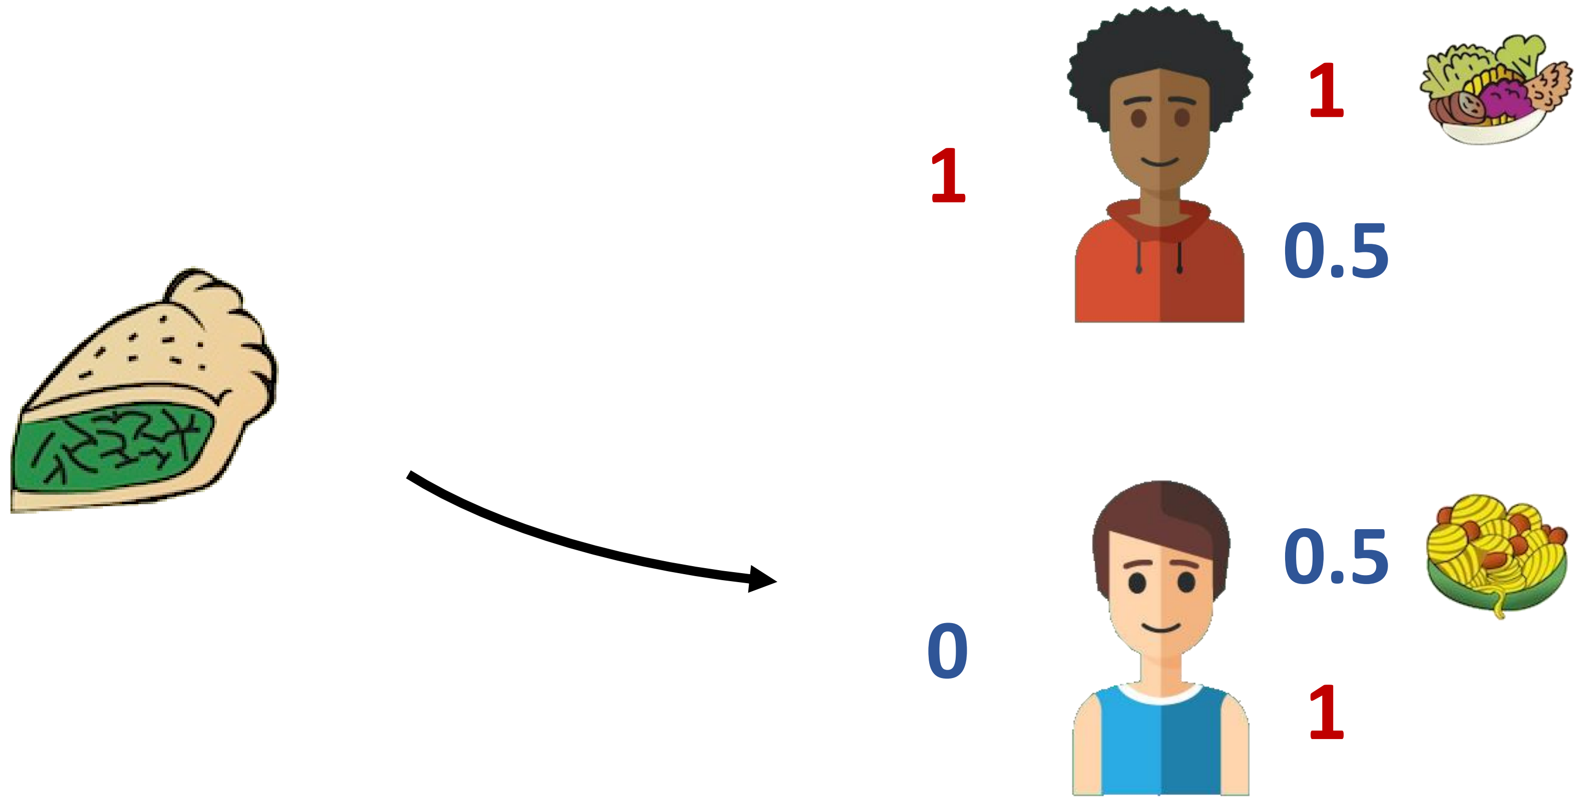
1
1
0.5
0.5
0
1

$$ENVY^{ij} = \max\{v_i(A_j) - v_j(A_i), 0\}$$

$$ENVY_3^{RB} = 1.5 - 1 = 0.5$$

$$ENVY_3^{BR} = 1 - 0.5 = 0.5$$

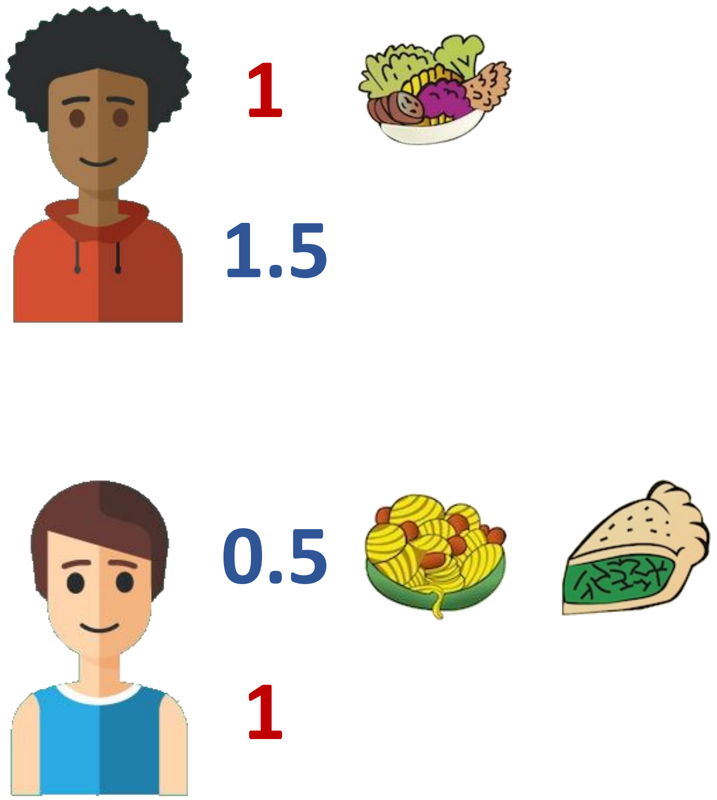

Model

- Items arrive online in batches of size k
- n agents
- Agent i has value $v_{it} \in [0,1]$ for item t
- Allocate to minimize $\mathrm{E}[ENVY_T]$
- Item values are chosen by adaptive adversary (maximizes $\mathrm{E}[ENVY_T]$)

Can we ensure *vanishing envy* after T items?

$$\left(\lim_{T\to\infty} \frac{ENVY_T}{T} = 0 \right)$$

Can we ensure *vanishing envy* at time T?

T batches of size 1

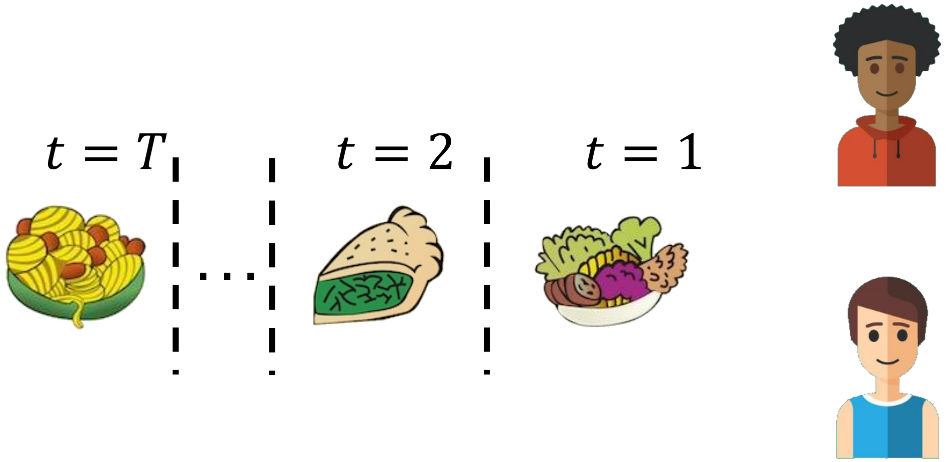

T/k batches of size k

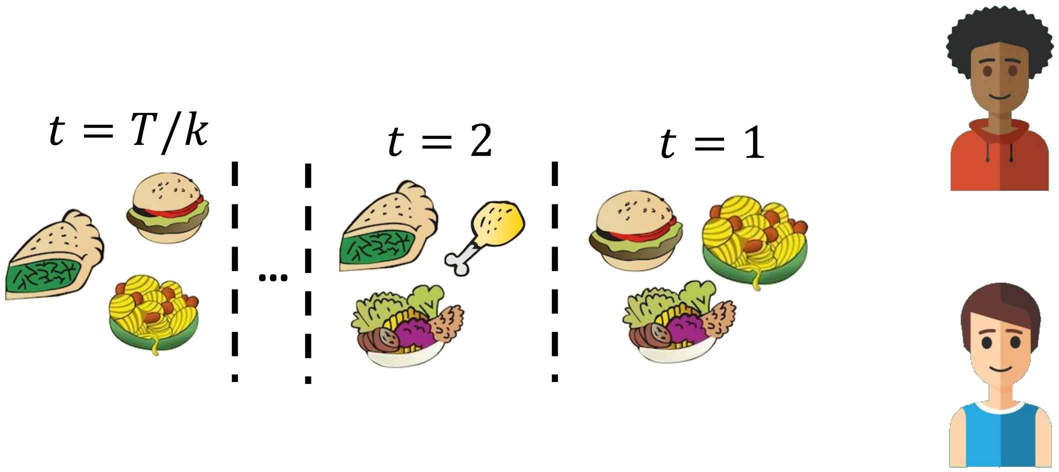

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$^1/_2$	1	ϵ	1	ϵ	…
Value agent 2	$^1/_2$	$^1/_4$	1	ϵ	1	…
Envy agent 1						
Envy agent 2						

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$^1/_2$	1	ϵ	1	ϵ	…
Value agent 2	$^1/_2$	$^1/_4$	1	ϵ	1	…
Envy agent 1						
Envy agent 2						

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$1/2$	1	ϵ	1	ϵ	…
Value agent 2	$1/2$	$1/4$	1	ϵ	1	…
Envy agent 1	$-1/2$					
Envy agent 2	$1/2$					

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$1/2$	1	ϵ	1	ϵ	…
Value agent 2	$1/2$	$1/4$	1	ϵ	1	…
Envy agent 1	$-1/2$	$1/2$				
Envy agent 2	$1/2$	$1/4$				

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$^1/_2$	1	ϵ	1	ϵ	…
Value agent 2	$^1/_2$	$^1/_4$	1	ϵ	1	…
Envy agent 1	$-^1/_2$	$^1/_2$	$^1/_2 - \epsilon$			
Envy agent 2	$^1/_2$	$^1/_4$	$^5/_4$			

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$1/2$	1	ϵ	1	ϵ	…
Value agent 2	$1/2$	$1/4$	1	ϵ	1	…
Envy agent 1	$-1/2$	$1/2$	$1/2 - \epsilon$	$3/2 - \epsilon$		
Envy agent 2	$1/2$	$1/4$	$5/4$	$5/4 - \epsilon$		

Greedy deterministic allocation

Give to whoever has the highest envy

Value agent 1	$1/2$	1	ϵ	1	ϵ	…
Value agent 2	$1/2$	$1/4$	1	ϵ	1	…
Envy agent 1	$-1/2$	$1/2$	$1/2-\epsilon$	$3/2-\epsilon$	$3/2-2\epsilon$	…
Envy agent 2	$1/2$	$1/4$	$5/4$	$5/4-\epsilon$	$9/4-\epsilon$	…

$$ENVY_T \approx T\ /\ 2$$

Items arrive one at a time

- Pretend that every agent values every item at 1.
- Allocate every item uniformly at random.

- Max change in $ENVY_{ij}$ is 1. T items arrive, i gets it w.p. $1/n$
- Random walk with step size 1, T/n steps. Deviation $\sim\sqrt{T/n}$

Theorem. Allocating unit-valued items uniformly at random yields $ENVY_T \in O\left(\sqrt{T \log n/n}\right)$ with high probability, and $\boldsymbol{E}[ENVY_T] \in O\left(\sqrt{T \log T/n}\right)$.

Items arrive in batches

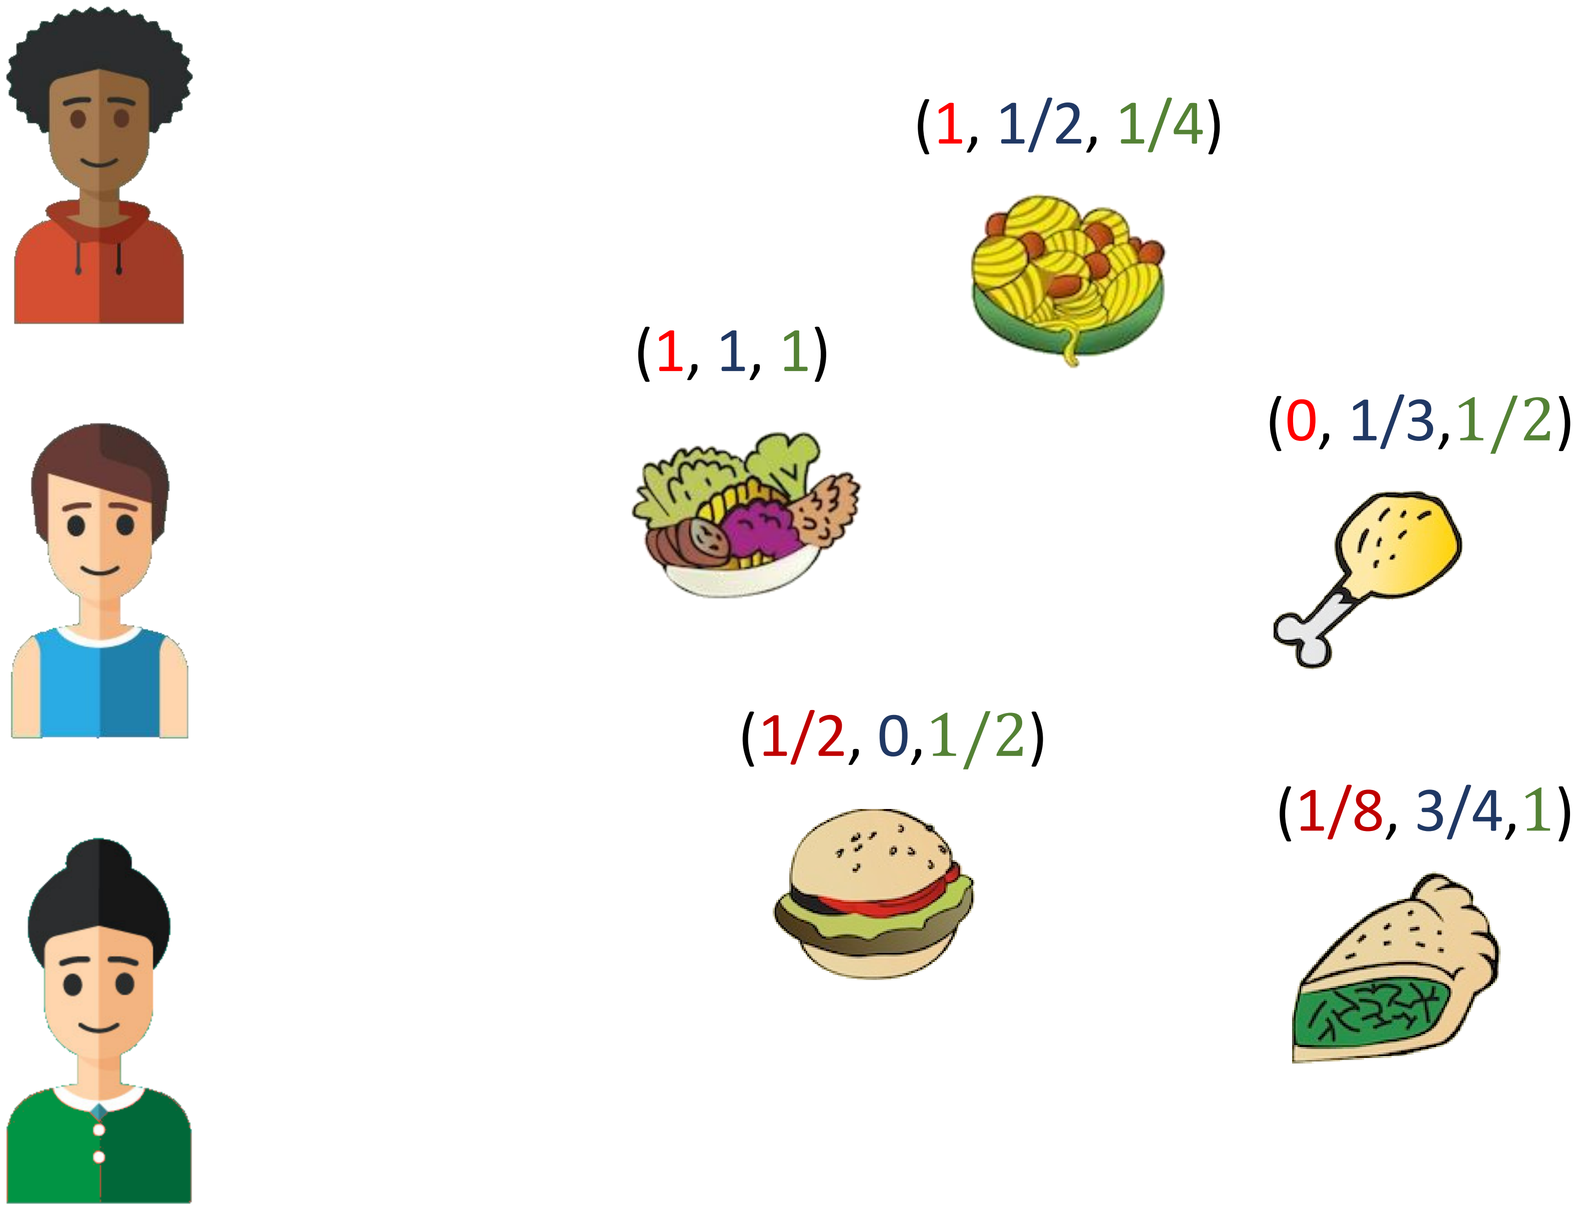
(1, 1/2, 1/4)
(1, 1, 1)
(0, 1/3,1/2)
(1/2, 0,1/2)
(1/8, 3/4,1)

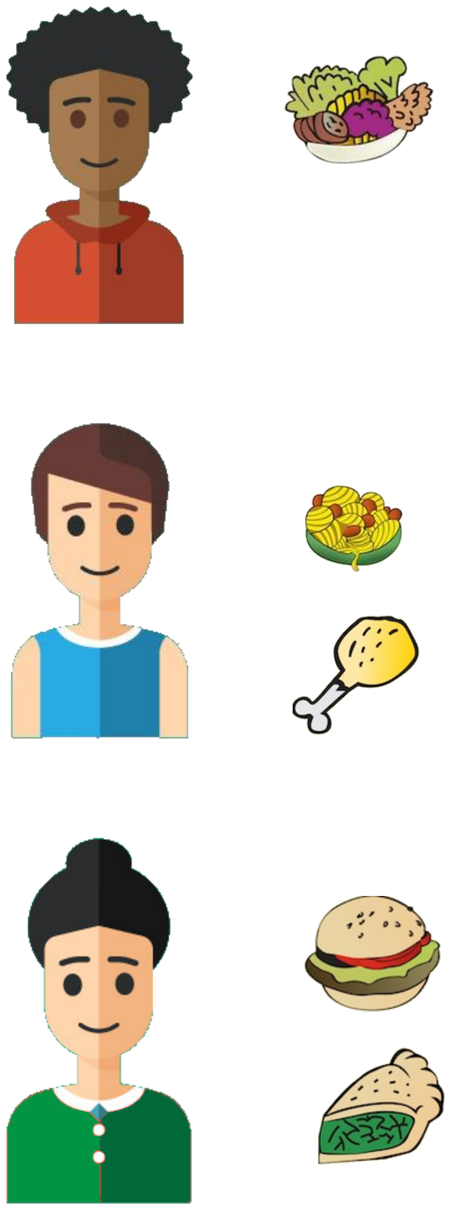

Next batch arrives

(T/k batches)

Intuition

- When $k = 1, ENVY_T \in \widetilde{\Theta}(\sqrt{T\ /\ n})$
- When $k = T, ENVY_T \leq 1$

$k = 1$

- Envy changed by ≤ 1 per round
- $ENVY^{ij}$ changed $\approx T\ /\ n$ times
- Bound: $ENVY_T \in \widetilde{\Theta}(1 \cdot \sqrt{T\ /\ n})$

General k

- Envy changes by ≤ 1 per round
- $ENVY^{ij}$ may change $T\ /\ k$ times
- Bound: $ENVY_T \in \widetilde{\Theta}(\sqrt{T\ /\ k})$?

Theorem. $ENVY_{T,k} \in \Omega(\sqrt{T/kn})$ and $ENVY_{T,k} \in \tilde{O}(\sqrt{T/k})$.

Upper bound (batches)

General k

- Envy changes by ≤ 1 per round
- $ENVY^{ij}$ may change $T \;/\; k$ times
- Bound: $ENVY_T \in \tilde{\Theta}(\sqrt{T \;/\; k})$?

Problems

- Can we find `balanced’ allocations across batches?

We don’t need EF1 (step size 1) per batch, any bounded constant change in envies will do.

Let’s try rounding near-integral envy-free allocations.

Finding near integral envy-free solutions

- x_{ij} fraction of item j given to agent i
- $v_{ij} \leq 1$ utility of agent i for item j

$$\sum_j v_{ij} x_{ij} \geq \sum_j v_{ij} x_{kj} \quad \forall i, k \quad \text{Envy-freeness}$$

$$\sum_i x_{ij} = 1 \quad \forall j \quad \text{Every item assigned}$$

$$x_{ij} \geq 0 \quad \forall i, j \quad \text{Non-negativity}$$

How many fractional values?
Number of variables: nm
Number of const: $\mathrm{n}^2 - \mathrm{n} + m$
So $\leq n^2 - n + m$ pos. variables

(Basically) tight, so $\sim n$ fractional items per agent

Rounding introduces n envy

Feasible solution is envy-free. We have to round it to get allocation

Rounding introduces envy: up to 1 per rounded item per agent

Theorem (Stromquist, 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

Theorem (Stromquist, 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

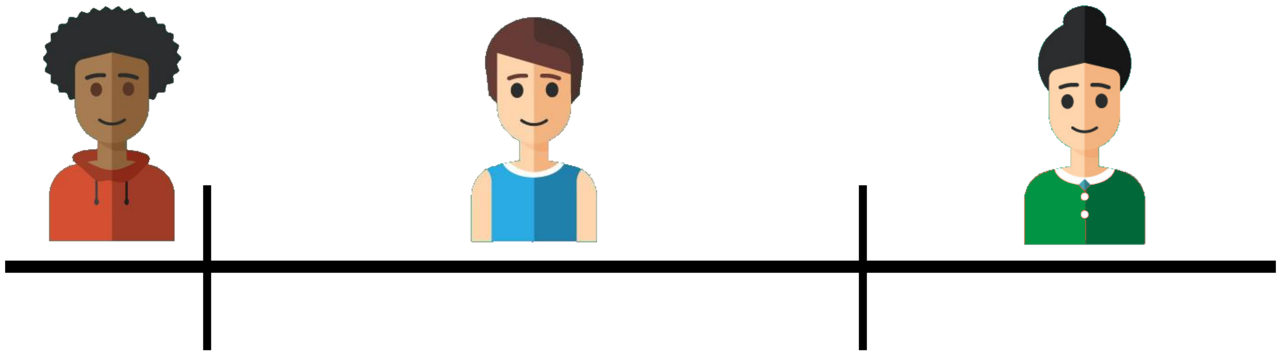

Theorem (Stromquist 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

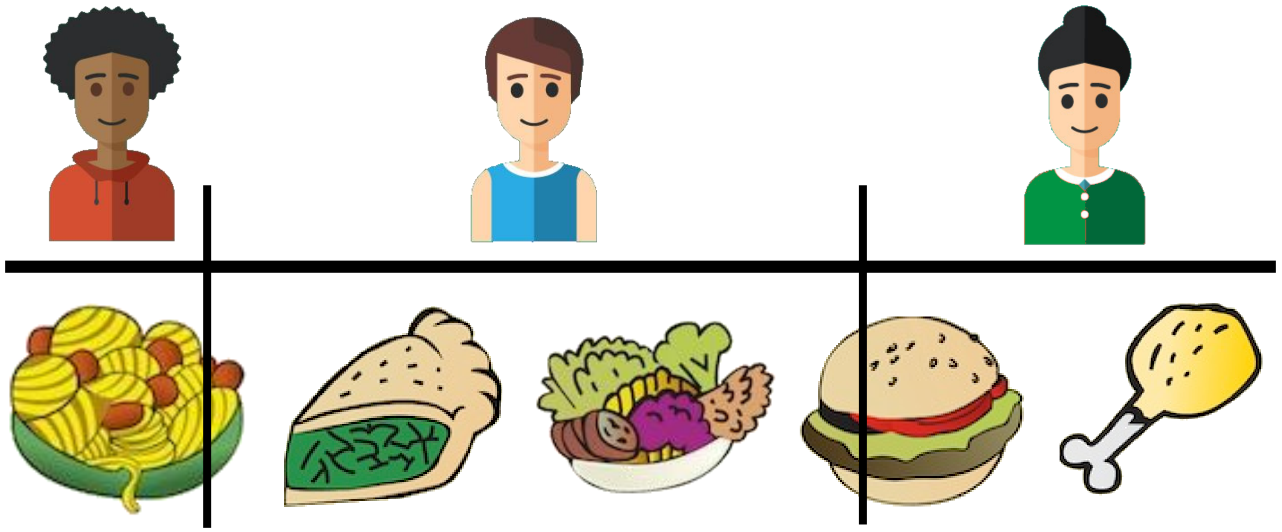

Theorem (Stromquist 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

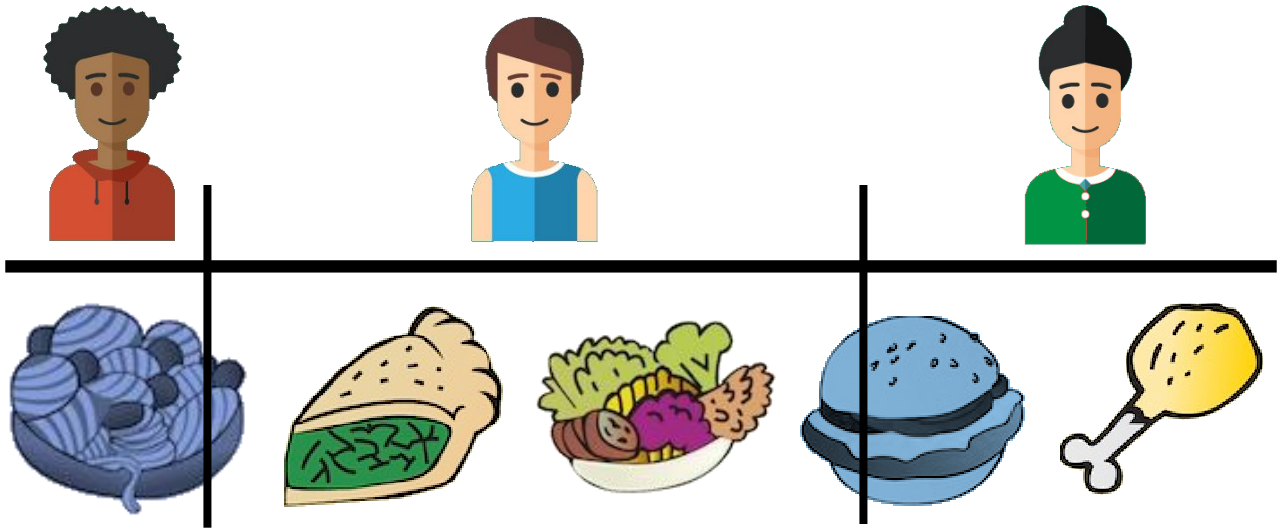

- No player receives more than two fractional items.

Theorem (Stromquist 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

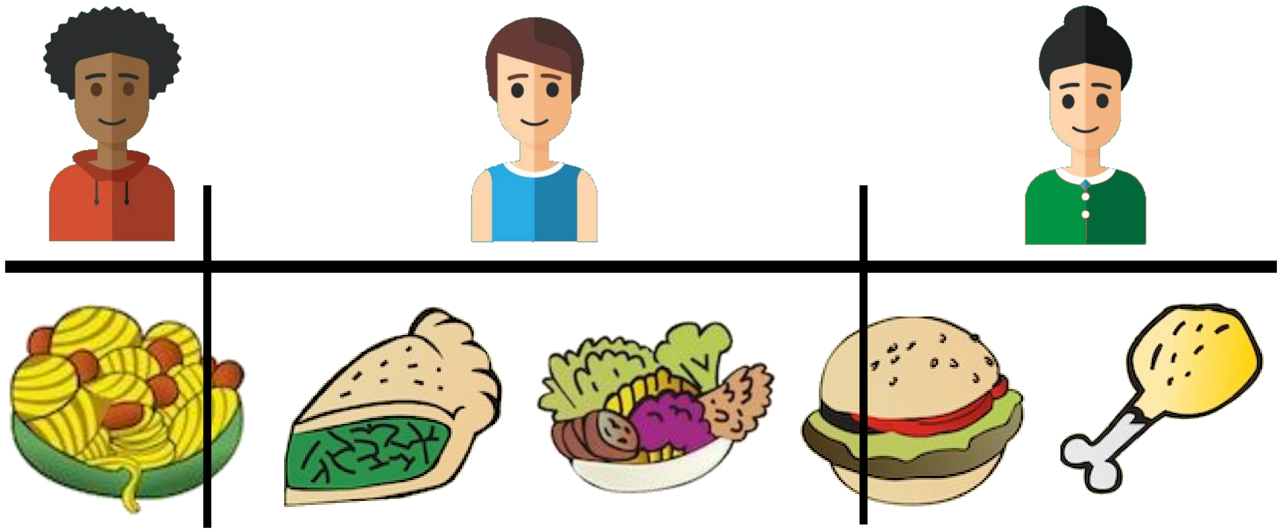

- No player receives more than two fractional items.
- Rounding fractional items randomly guarantees envy changes by ≤ 4

Theorem (Stromquist 1980). Suppose n agents have `reasonable' valuations over $[0,1]$. Then there exists an envy-free division of the interval in which every agent receives a single contiguous segment.

- x_{ij} fraction of item j given to agent i
- $v_{ij} \leq 1$ utility of agent i for item j
- Indicator variables x_{ij}^0, x_{ij}^1: sum to 0 when x_{ij} is fractional, sum to 1 o.w.

$\sum_j v_{ij} x_{ij} \geq \sum_j v_{ij} x_{kj} \;\; \forall i, k$	Envy-freeness
$\sum_i x_{ij} = 1 \;\; \forall j$	Every item assigned
$x_{ij} \geq 0 \;\; \forall i, j$	Non-negativity
$x_{ij}^0 \leq x_{ij} \leq 1 - x_{ij}^1 \;\; \forall i, j$	Indicator variable constraints
$\sum_j \left(x_{ij}^0 + x_{ij}^1 \right) \geq m - 2 \;\; \forall i$	At most 2 fractional items per agent
$x_{ij}^0, x_{ij}^1 \in \{0,1\} \;\; \forall i, j$	

Questions