Fairness in Non-Truthful Algorithms with Strategic Agents

Talk at the Amsterdam/Saint-Etienne Workshop on Social Choice,
March 14, 2024 – Rebecca Reiffenhäuser

Based largely on joint work with
Georgios Amanatidis, Georgios Birmpas, Philip Lazos, Stefano Leonardi:

- Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness – WINE 2021 (best paper), and
- Round-Robin Beyond Additive Agents: Existence and Fairness of Approximate Equilibria – EC 2023

Introduction and Overview

Fair Allocation

Fairness is a natural goal in:

Divorce Settlement, Inheritance, Cost Sharing in Communication Networks, Distributed Resource Allocation/Wireless Systems, Birthday Parties (Cake Cutting)..

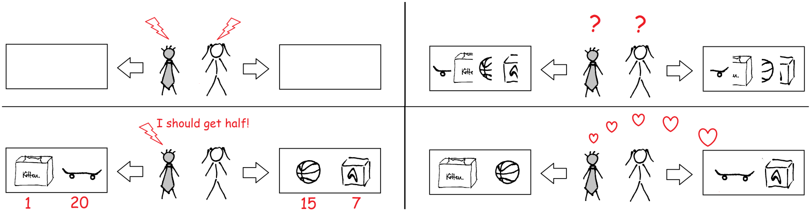

Main Problems:

What's 'fair'? People have subjective ideas of that...

Can we be fair? Often, not really...

Fairness is a natural goal in:

Divorce Settlement, Inheritance, Cost Sharing in Communication Networks, Distributed Resource Allocation/Wireless Systems, Birthday Parties (Cake Cutting)..

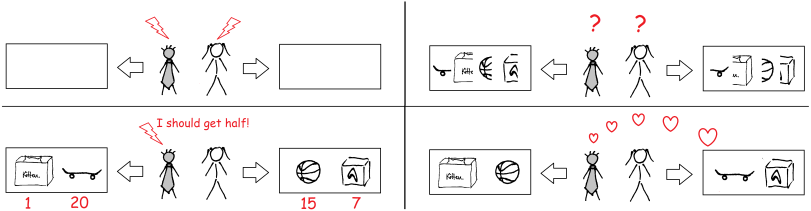

Assumptions:

No Free Disposal Can't just throw away items to make the outcome 'fair'.

Indivisible Goods Can't just give everyone a 'fair share' of each good!

Formal Problem Definition

Additive Combinatorial Assignment

Given a set M of m indivisible goods, and a set N of n agents. Each agent has an additive valuation function defining

$$v_{ij} = \text{ value derived by agent } i \text{ for obtaining good } j$$

and $v_i(S) = \sum_{j \in S} v_{ij}$, for any set S of goods.

An *allocation* is a partition $S = (S_1, S_2, ..., S_n)$ of the set M of goods.

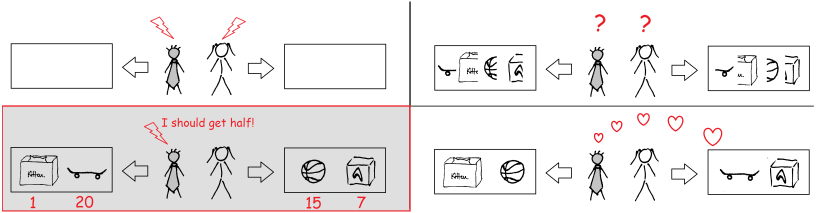

Proportionality (PROP) [Steinhaus, 1949]

An allocation $S = \{S_1, S_2, \ldots S_n\}$ is *proportional* if each agent receives at least $1/n$ of his total valuation on all goods, where $n = |N|$.

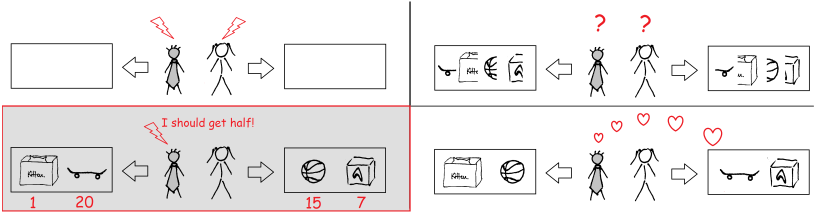

Proportionality (PROP) [Steinhaus, 1949]

An allocation $S = \{S_1, S_2, \dots S_n\}$ is *proportional* if each agent receives at least $1/n$ of his total valuation on all goods, where $n = |N|$.

Maximin-Share Fairness (MMS) [Budish, 2011]

An allocation $S = \{S_1, S_2, \dots S_n\}$ is *MMS* if each agent receives at least the value he would when after optimally dividing all goods into n bundles, the *worst* of these is assigned to him.

Fairness Notions

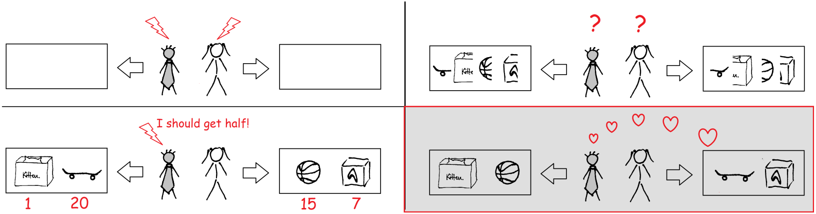

Envy-Freeness (EF) [Foley '67, Varian '74]

An allocation $S = \{S_1, S_2, \ldots S_n\}$ is envy-free if for all $i, j \in N$

$$v_i(S_i) \geq v_i(S_j)$$

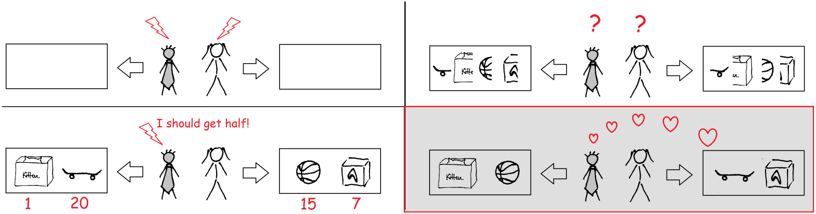

Envy-Freeness (EF) [Foley '67, Varian '74]

An allocation $S = \{S_1, S_2, \ldots S_n\}$ is envy-free if for all $i, j \in N$

$$v_i(S_i) \geq v_i(S_j)$$

Envy-Fr. up to Any Good (EFX) [Gourves et al. '14, Carag. et al. '19]

An allocation $S = \{S_1, S_2, \ldots S_n\}$ is EFX if for all $i, j \in N$

$$v_i(S_i) \geq v_i(S_j \setminus \{m_{min}\}),\text{, where } m_{min} = \min_{m \in S_j} \{v_{im}\}$$

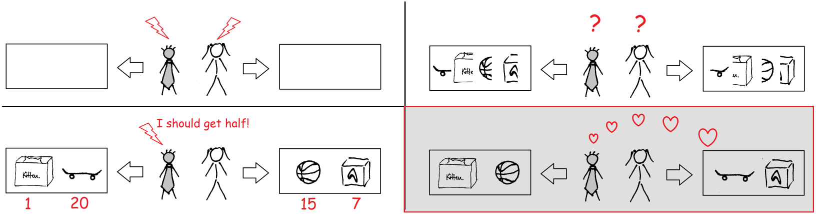

Envy-Freeness (EF) [Foley '67, Varian '74]

An allocation $S = \{S_1, S_2, \dots S_n\}$ is envy-free if for all $i, j \in N$

$$v_i(S_i) \geq v_i(S_j)$$

Envy-Fr. up to One Good (EF1) [Lipton et al. '04, Budish '11]

An allocation $S = \{S_1, S_2, \dots S_n\}$ is EF1 if for all $i, j \in N$

$$v_i(S_i) \geq v_i(S_j \setminus \{m_{max}\}), \text{, where } m_{max} = \max_{m \in S_j} \{v_{im}\}$$

Goals: Fairness and Incentive Compatibility

Fairness Notion

Guarantee, e.g.

- Proportionality (PROP)
- Maximin-Share Fairness (MMS)
- Envy-Freeness (EF)
- Envy-Freeness up to Any Good (EFX)
- **Envy-Freeness up to One Good (EF1)**

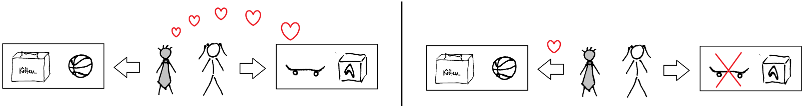

Incentive-Compatibility/Truthfulness

Give a mechanism (here: an algorithm, but with inputs that might misrepresent the actual values) such that it is every agent's best strategy to report their true valuations.

Goals: Fairness and Incentive Compatibility

Fairness Notion

Guarantee, e.g.

- Proportionality (PROP)
- Maximin-Share Fairness (MMS)
- Envy-Freeness (EF)
- Envy-Freeness up to Any Good (EFX)
- **Envy-Freeness up to One Good (EF1)**

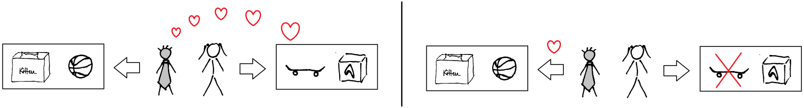

Incentive-Compatibility/Truthfulness

Give a mechanism (here: an algorithm, but with inputs that might misrepresent the actual values) such that it is every agent's best strategy to report their true valuations.

Can we combine **fair and truthful** in the same routine?

Impossibility: Fairness and Incentive Compatibility

Can we combine **fair and truthful** in the same routine?

Impossibility: Fairness and Incentive Compatibility

Can we combine **fair and truthful** in the same routine?

Impossibility Results [Amanatidis et al. 2017]

Truthfulness and Fairness are **incompatible** even for only two players, for various fairness notions!

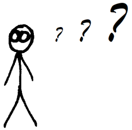

Impossibility: Fairness and Incentive Compatibility

Can we combine **fair and truthful** in the same routine?

Impossibility Results [Amanatidis et al. 2017]

Truthfulness and Fairness are **incompatible** even for only two players, for various fairness notions!

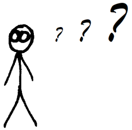

Modified Goal:

Can we at least have *non-truthful* mechanisms,
but which have equilibria that define fair allocations?

Goals: Fair, Truthful Equilibria

Envy-Freeness up to One Good (EF1)

Produce an allocation such that everyone would prefer their own, assigned items over the set of anyone else - after the *best* item from the other person's set was taken out...

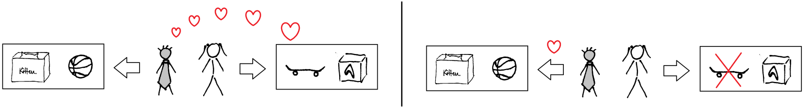

'Truthful Equilibria'

Give a mechanism (here: an algorithm, but with inputs that might misrepresent the actual values) such that the produced outcome is an equilibrium (PNE) with respect to the *actual* valuations?

Note: by saying Equlibrium here, we mean Pure Nash Equilibrium (PNE), i.e.: when fixing the reports of all others, no agent can obtain a better assignment by modifying his own report.

The Algorithm: Round-Robin

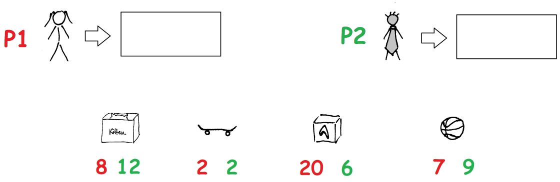

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

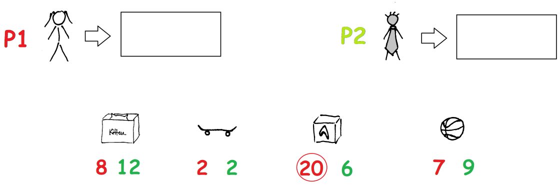

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

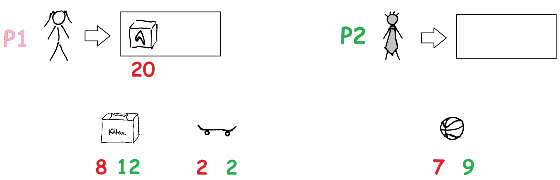

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

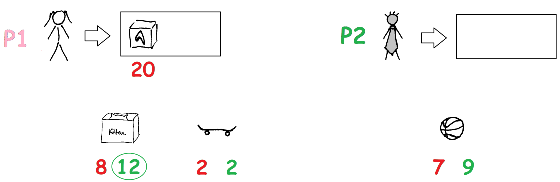

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

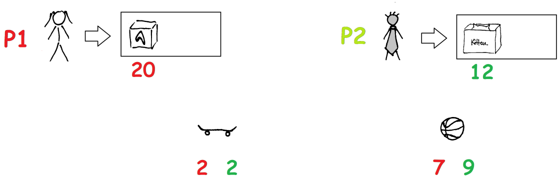

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

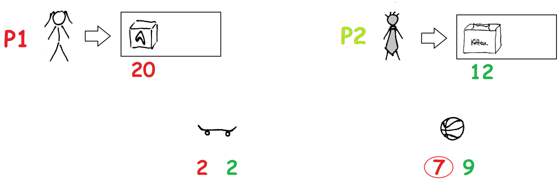

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

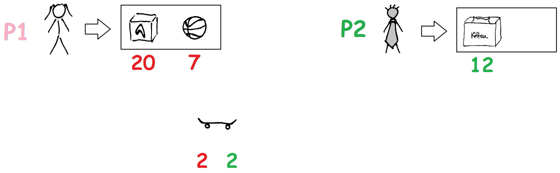

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

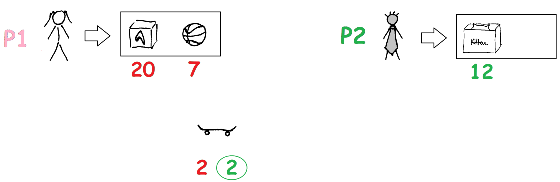

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

The Algorithm: Round-Robin

Players $1, 2, \ldots, n$ take turns being allocated their *one* highest-valued good from the available ones, according to this order, until all goods are allocated.

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

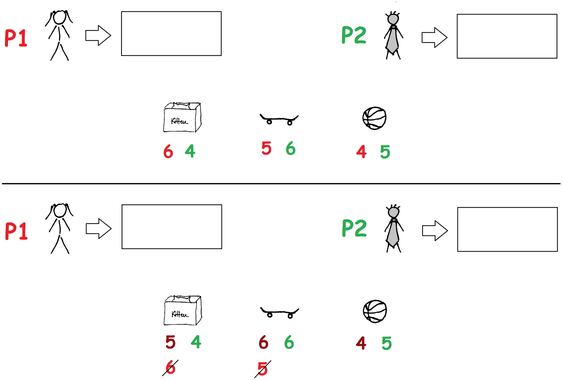

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

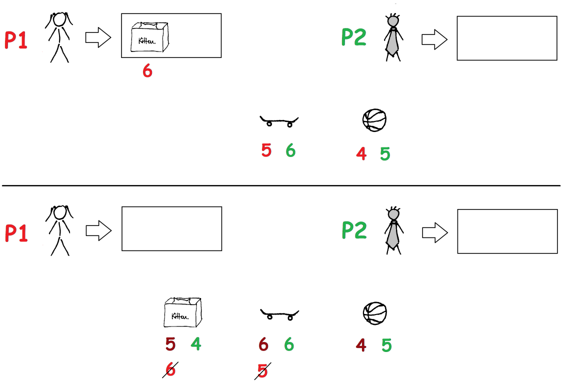

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

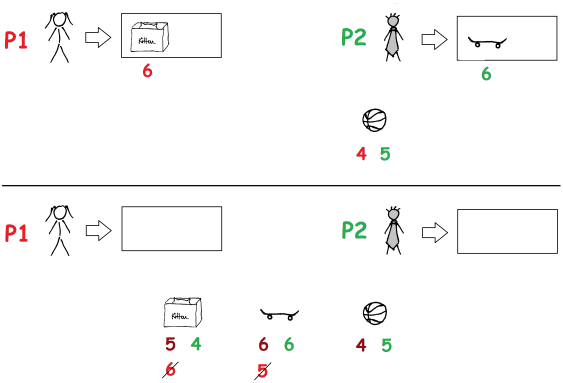

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

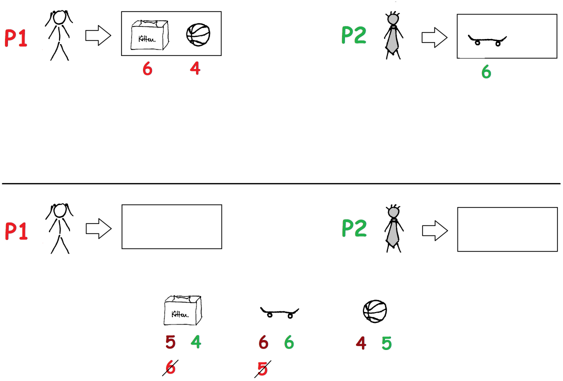

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

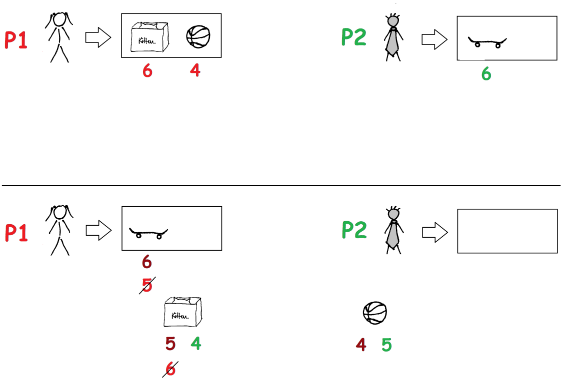

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

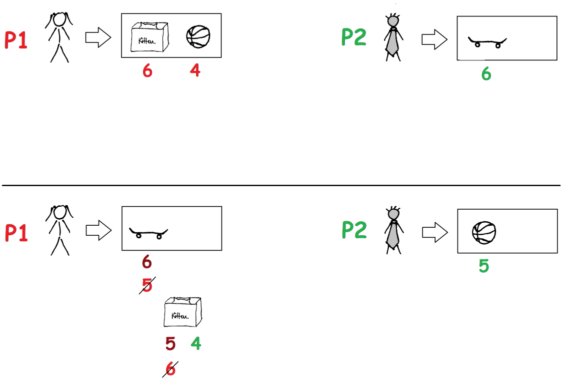

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

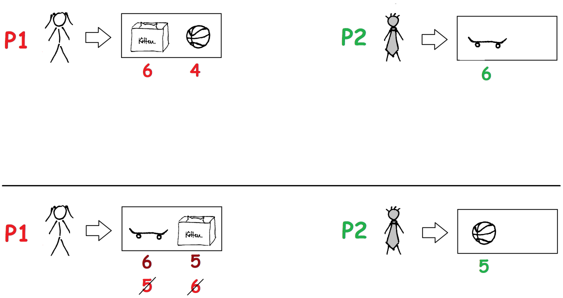

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

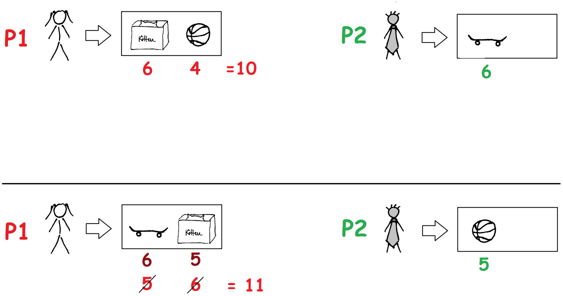

Example: Truthful Reports don't make Equilibria

Saying the truth, in general, does *not* lead to a PNE, see the following minimal example with two agents and three goods:

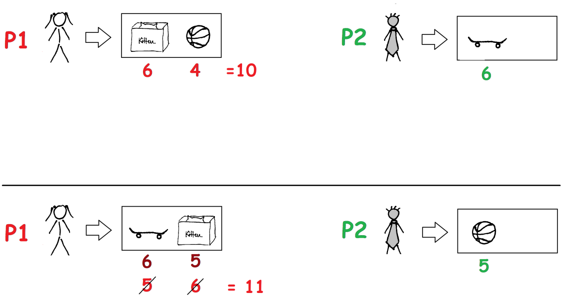

truth-telling is NOT a PNE!

Key Observation: EF1 for all, EF for Agent 1

Round-Robin Algorithm produces EF1 Allocations

When presented with the true, additive valuations of the agents, all outcomes of Round-Robin are EF1 **[Markakis 2017, Caragiannis et al. 2019]**.

- **Round-Robin (as an algorithm, on the true values) is EF from view of agent** 1
  since, in every round, he gets more than anyone else!
- **Round-Robin is EF1 from view of any agent** i
  since after ignoring all previous $i-1$ picks of the first round, we can just pretend he's agent one!

Key Observation: EF1 for all, EF for Agent 1

Round-Robin Algorithm produces EF1 Allocations

When presented with the true, additive valuations of the agents, all outcomes of Round-Robin are EF1 **[Markakis 2017, Caragiannis et al. 2019]**.

- **Round-Robin (as an algorithm, on the true values) is EF from view of agent** 1
  since, in every round, he gets more than anyone else!
- **Round-Robin is EF1 from view of any agent** i
  since after ignoring all previous $i-1$ picks of the first round, we can just pretend he's agent one!

Additive Valuations

Round-Robin Results: Additive Valuations

Round-Robin always produces EF1 Allocations

When presented with the true, additive valuations of the agents, all outcomes of Round-Robin are EF1 **[Markakis 2017, Caragiannis et al. 2019]**.

Round-Robin always has PNE

The Bluff Profile is always a PNE of Round-Robin for additive valuations. **[Aziz et al. 2017]**

We showed:

Round-Robin Mechanism produces truthful EF1-PNE!

For every instance $\mathcal{I} = \{N, M, v\}$, each PNE of Round-Robin is EF1 with respect to the true valuations.
[Amanatidis, Birmpas, Lazos, Leonardi, R.R. 2021]

PNE-Alloc.

P1	P2	P3	P1	P2	P3	P1	P2	
22	2	20	15	11	6	3		b_1
10	6	11	8	9	4	3	<3	v_1

Main Technical Lemma:

Assume b_1 is a best response of agent 1 to $b_{-1} = (b_2, b_3, \ldots, b_n)$. Then, there exists a valuation function b_1^* such that:

- Round-Robin produces the same allocation $(S_1, \ldots, S_n)$ on b as on (b_1^*, b_{-1}).
- $b_1^*(S_1) = v_1(S_1)$
- $b_1^*(j) = v_1(j)$ for all $j \in M \setminus S_1$.

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

PNE-Alloc.

	P1	P2	P3	P1	P2	P3	P1	P2
b_1	22	2	20	15	11	6	3	
v_1	10	6	11	8	9	4	3	<3

Main Technical Lemma:

Assume b_1 is a best response of agent 1 to $b_{-1} = (b_2, b_3, \ldots, b_n)$. Then, there exists a valuation function b_1^* such that:

- Round-Robin produces the same allocation $(S_1, \ldots, S_n)$ on b as on (b_1^*, b_{-1}).
- $b_1^*(S_1) = v_1(S_1)$
- $b_1^*(j) = v_1(j)$ for all $j \in M \setminus S_1$.

Note: The allocation $(S_1, \ldots, S_n)$ is then EF with regards to b_1^* as well as v_1!

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

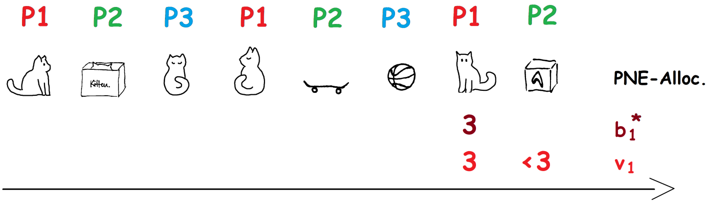

Intuitive Strategy:

It is not harmful to agent 1 if he plays truthfully in the very last round.

However: Does not hold inductively for previous rounds (see ex.)!

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

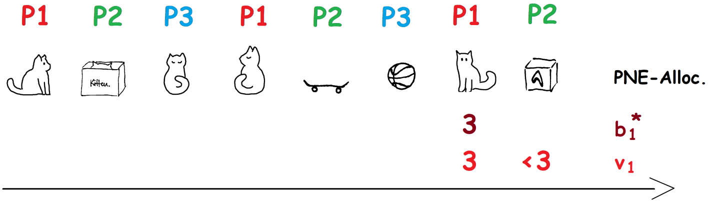

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

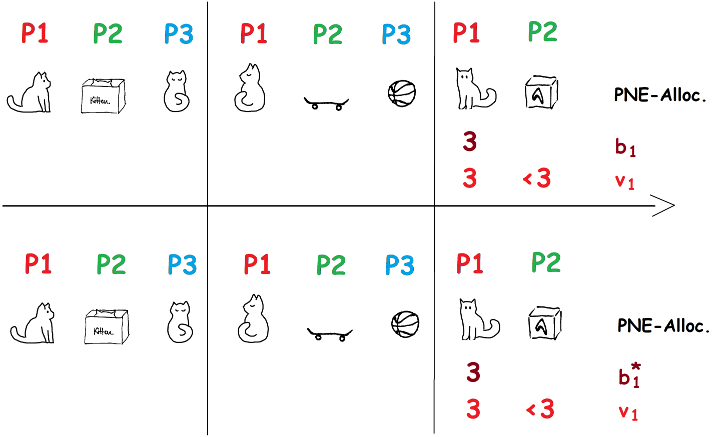

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

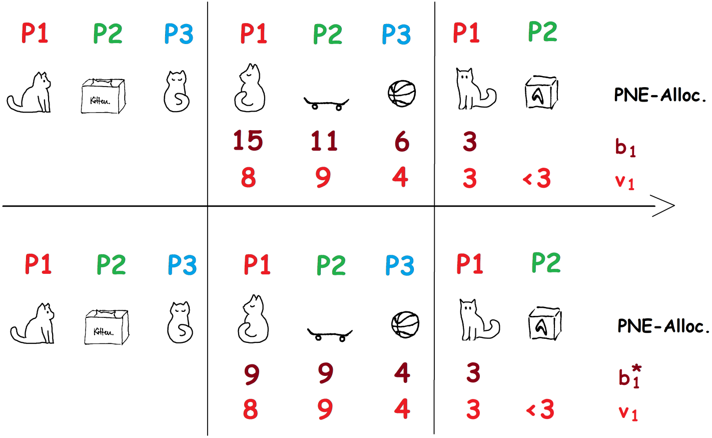

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

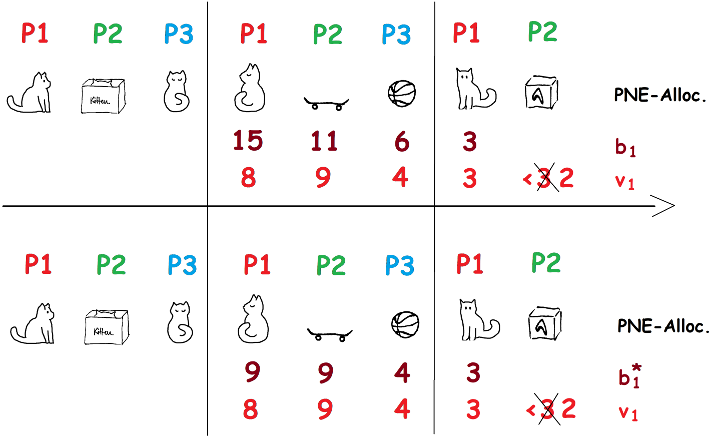

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

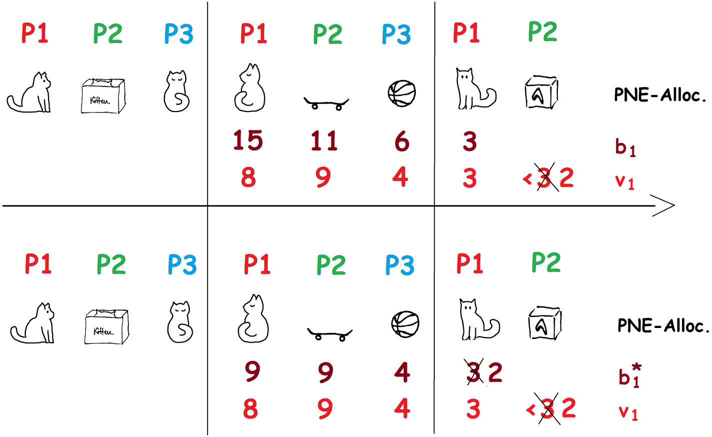

Proof Sketch: Round-Robin PNE are Truthful EF1 for Additive

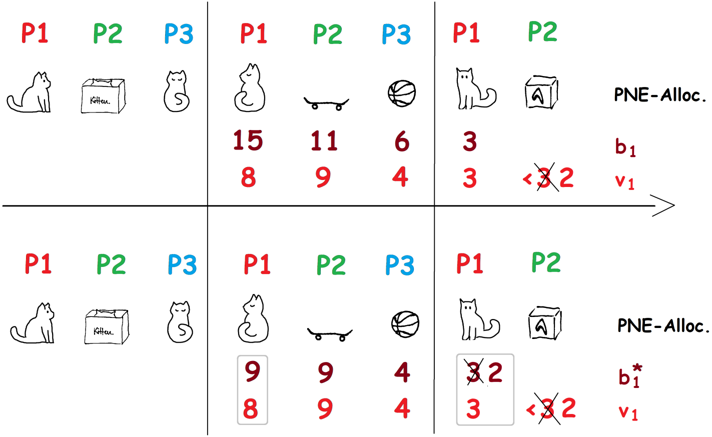

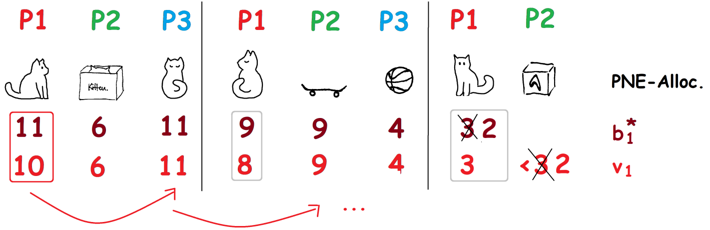

- Replace values in b_1 in order-preserving way (no changes to alloc.)
- Observe: When $\sum_{j \in S_1} v_{1j} \neq \sum_{j \in S_1} b_1^*(j)$ because the alloc. would change: **shift along a <u>chain</u> of adjustments**!
  (Possible, since set of available items would only change by one)
- Lucky: Because of essentially the EF-property of Round-Robin, value adjustments always work out!

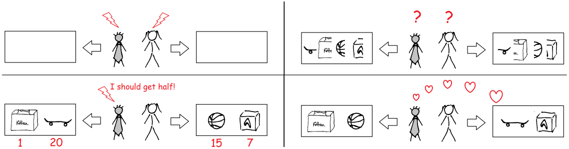

Round-Robin for additive valuations

– always has pure Nash equilibria and

– these induce allocations that are EF1 w.r.t. the underlying true values.

That is, Round-Robin retains its fairness properties at its equilibria, even when the input is given by strategic agents!

Beyond Additive Valuations

Moving On: Beyond Additive Valuations

For each agent $i \in N$, we say that v_i is

- **subadditive**, if $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every $S, T \subseteq M$.
- **submodular**, if $v_i(g \mid S) \geq v_i(g \mid T)$ for any $S \subseteq T \subseteq M$ and $g \notin T$.
- **cancelable**, if $v_i(S \cup \{g\}) > v_i(T \cup \{g\}) \Rightarrow v_i(S) > v_i(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.
- **additive**, if $v_i(S \cup T) = v_i(S) + v_i(T)$ for every $S, T \subseteq M$ with $S \cap T = \emptyset$.

Beyond Additive: What's Different

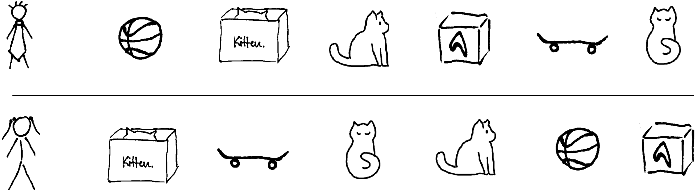

1. **First, Round-Robin needs to be clearly defined**
   (what is the next-best item when v is some set function?)

Beyond Additive: What's Different

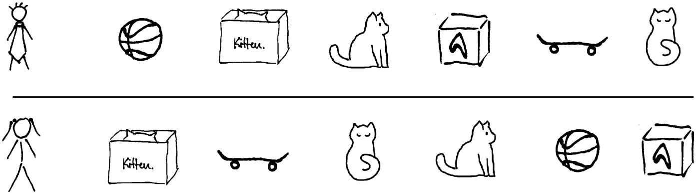

1. **First, Round-Robin needs to be clearly defined (what is the next-best item when v is some set function?)**
2. **Do PNE still exist?**

Beyond Additive: What's Different

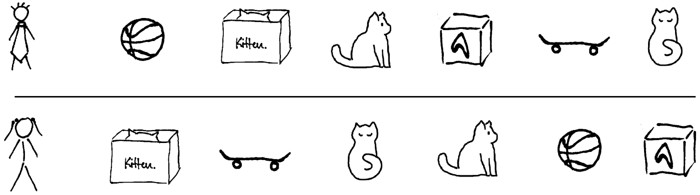

1. **First, Round-Robin needs to be clearly defined**
   (what is the next-best item when v is some set function?)

2. **Do PNE still exist?**

3. **If so, is there still a truthful (EF1) fairness guarantee?**

Beyond Additive: What's Different

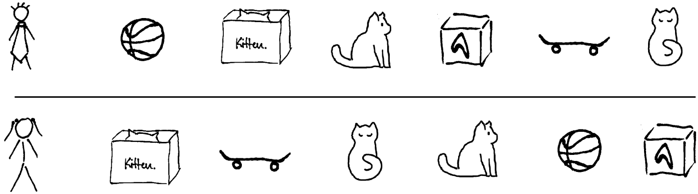

1. **First, Round-Robin needs to be clearly defined (what is the next-best item when v is some set function?)**
   We simply define it via assigning the next item due to each agent's submitted *preference order*.
2. **Do PNE still exist?**
3. **If so, is there still a truthful (EF1) fairness guarantee?**

Beyond Additive: What's Different

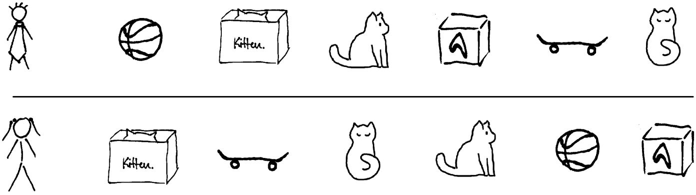

1. **First, Round-Robin needs to be clearly defined (what is the next-best item when v is some set function?)**
   We simply define it via assigning the next item due to each agent's submitted *preference order*.
2. **Do PNE still exist?**
   Surprisingly, largely yes (at least approximately), and we can still construct them!
3. **If so, is there still a truthful (EF1) fairness guarantee?**

Beyond Additive: What's Different

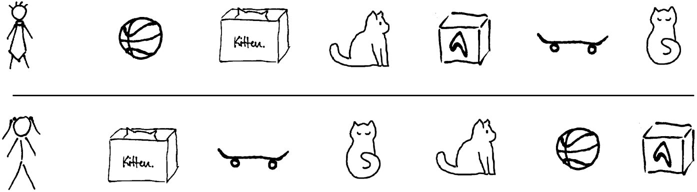

1. **First, Round-Robin needs to be clearly defined (what is the next-best item when v is some set function?)**
   We simply define it via assigning the next item due to each agent's submitted *preference order*.
2. **Do PNE still exist?**
   Surprisingly, largely yes (at least approximately), and we can still construct them!
3. **If so, is there still a truthful (EF1) fairness guarantee?**
   Surprisingly, yes! (At least approximately)

Cancelable Valuations

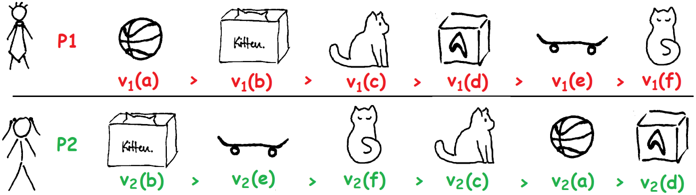

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

Cancelable Valuations

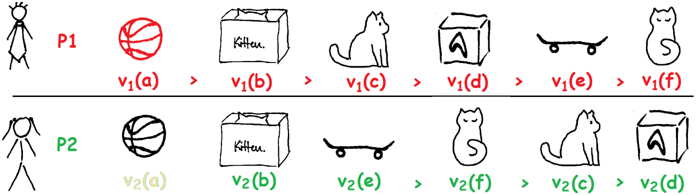

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $\mathit{argmax}_{g \in T} v(g) \subseteq \mathit{argmax}_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

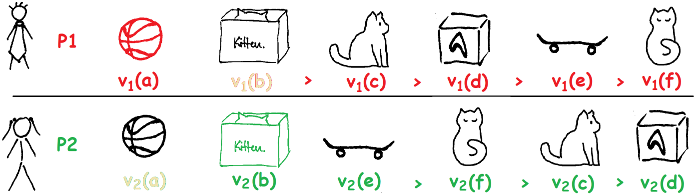

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

Cancelable Valuations

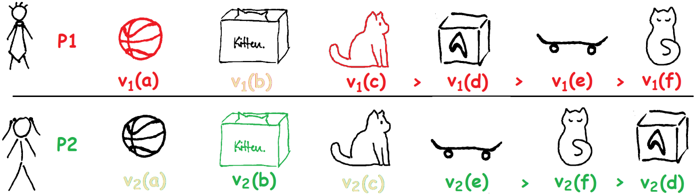

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

Cancelable Valuations

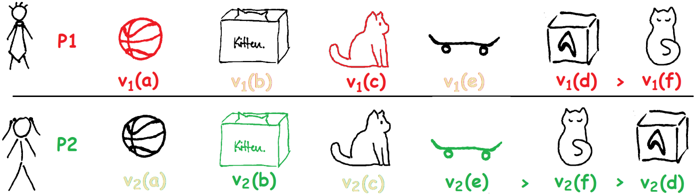

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

Cancelable Valuations

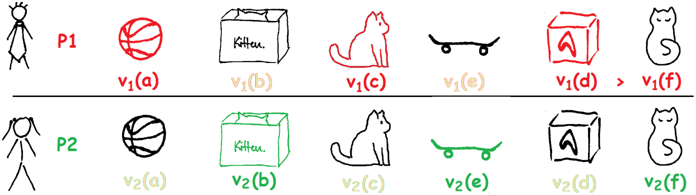

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

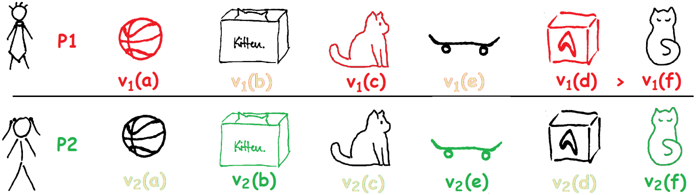

Reminder: valuation v is *cancelable*, if $v(S \cup \{g\}) > v(T \cup \{g\}) \Rightarrow v(S) > v(T)$ for any $S, T \subseteq M$ and $g \in M \setminus (S \cup T)$.

Useful Implication: $argmax_{g \in T} v(g) \subseteq argmax_{g \in T} v(g \mid S)$

Round-Robin and Bluff Profile for Cancelable Valuations

- Pick goods in RR according to best **singleton value** $v(g)$, $g \in M$.
- Define Bluff Profile via the resulting order, analogously to additive case.

Our Results: Cancelable Valuations

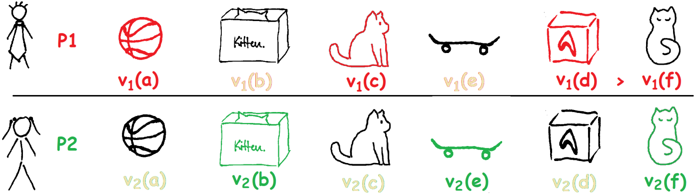

PNE via Bluff: Cancelable

When all agents have cancelable valuation functions,

- the Bluff Profile is an exact PNE of Round-Robin, and
- the obtained allocation is EF1.

Fairness Properties of General (Approx.) PNE: Cancelable

When all agents have subadditive cancelable valuation functions, any α-approx. PNE of Round-Robin results in an $\alpha/2$-EF1 allocation.

Submodular Valuations

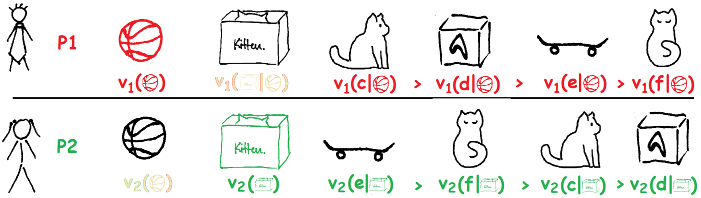

Reminder: valuation v is *submodular*, if $v_i(g \mid S) \geq v_i(g \mid T)$ for any $S \subseteq T \subseteq M$ and $g \notin T$.

Useful Implication: The *Greedy* Routine which assigns goods one-by-one, in each step realizing the maximum-possible marginal gain, is a 2-approximation.

Round-Robin and Bluff Profile for Submodular Valuations

- Pick goods in RR according to best available **marginal value** $v(g|S)$, $g \in M \setminus \{S\}, S \subseteq M$.
- Define Bluff Profile via the according assignment order.

Our Results: Submodular Valuations

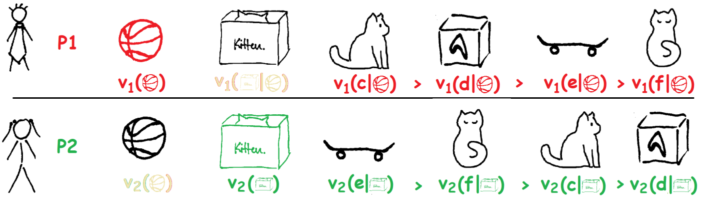

PNE via Bluff: Submodular

- The generalized Bluff Profile always is a $1/2$-PNE of Round-Robin, and this is tight (i.e., for any $\epsilon > 0$ there exist instances where it is not a $(1/2 + \epsilon)$-PNE).
- The allocation produced by Bluff is always $1/2$-EF1 with respect to the true valuations $v_i, i \in [N]$. This is tight (i.e., for any $\epsilon > 0$, there exist instances where this allocation is not $(1/2 + \epsilon)$-EF1.

Fairness Properties of General (Approx.) PNE: Submodular

When all agents have submodular valuation functions, any α-approx. PNE of Round-Robin results in an $\alpha/3$-EF1 allocation.

Overview: Round-Robin and its Truthful, Fair Equilibria

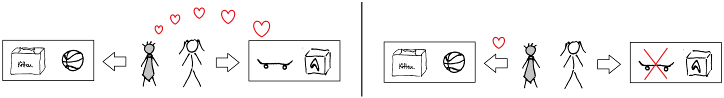

We showed, w.r.t. to the **true, private** valuations:

Additive Valuations:

RR always has PNE, e.g. the Bluff Profile, and all PNE are EF1.

Cancelable Valuations:

- RR always has PNE, e.g. Bluff by Singletons, and all PNE are EF1.
- For *subadditive cancelable*, any α-approx. PNE results in a $\alpha/2$-EF1 allocation.

Submodular Valuations:

- For some instances, no $(3/4 + \epsilon)$-PNE of RR exists.
- RR always has 1/2-PNE, e.g. Bluff by Marginals, which is always 1/2-EF1.
- Any α-approx. PNE results in an $\alpha/3$-EF1 allocation.

Future Directions

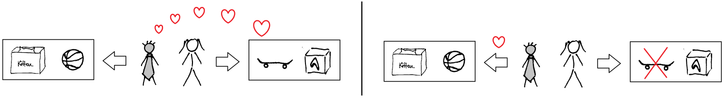

- Other, even more general valuation classes?
- Other algorithms?
- Other forms of alignment between selfish behavior and fairness goals?

The End.

Thank you!