On Existence of Truthful Fair Cake Cutting Mechanisms

Biaoshuai Tao (陶表帅)
John Hopcroft Center for Computer Science
Shanghai Jiao Tong University

Accepted in EC’22

Cake Cutting Problem

- How can we divide a divisible heterogeneous good (e.g., cake) fairly among agents?
- Divisible:
  - cake, time, land, computation resources, etc.
- Heterogeneous:
  - different agents have different preferences on different parts of the cake
- Fair:
  - Envy-free: no agent envies the allocation of any other agent
  - Proportional: each agent receives at least his/her average share

Applications

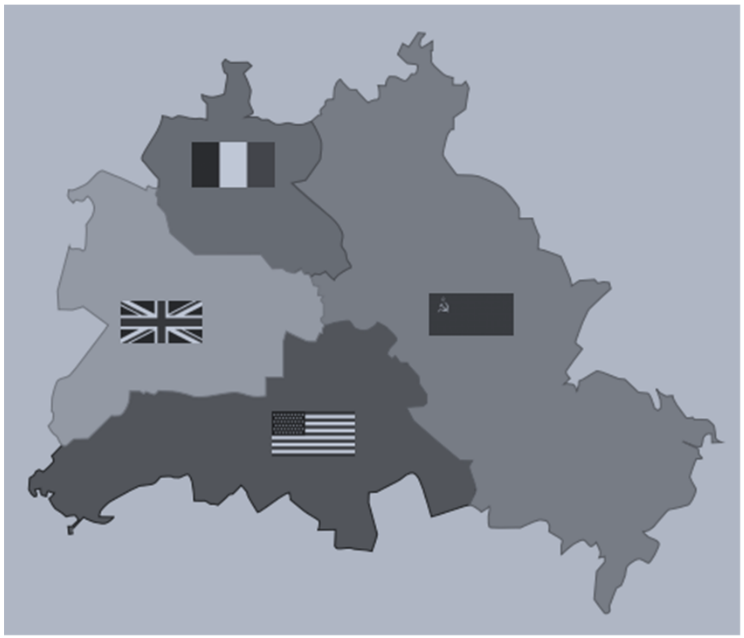

Model

- Cake: interval $[0, 1]$, allocated to n agents
- Allocation: $A = \{A_1, \dots, A_n\}$, a partition of $[0, 1]$
- Value density function $f_i: [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ for each agent i.
  - i's value on $S \subseteq [0, 1]$ is then

$$v_i(S) = \int_S f_i(x) dx$$

Fairness

- Envy-freeness: no one "envies" any other agent

$$\forall i, j: \ v_i\ (A_i) \geq v_i(A_j)$$

- Proportionality: each agent receives at least average value

$$\forall i: \ v_i\ (A_i) \geq \frac{1}{n} v_i([0, 1])$$

Relationship between Two Fairness Notions

- Envy-freeness implies proportionality.
- Envy-freeness: $v_i(A_i) \geq v_i(A_j)$ for $j = 1, \ldots, n$
- Summing up: $n \cdot v_i(A_i) \geq \sum_{j=1}^{n} v_i(A_j) = v_i(A_1 \cup \cdots \cup A_n) = v_i([0,1])$
- Implying proportionality: $v_i(A_i) \geq \frac{1}{n} v_i([0,1])$
- However, the implication fails if we are allowed to leave some part of the cake unallocated…
- Throwing away the entire cake (i.e., $A_1 = \cdots = A_n = \emptyset$) is envy-free, but not proportional.

Representation of f_i

- How to represent value density function succinctly?
- Approach 1: Direct Revelation: Assume each f_i is piecewise constant:
  - $\{X_1, \dots, X_m\}$: partition of $[0, 1]$, each f_i is constant on each X_t
- Approach 2: Robertson-Webb query model [Robertson & Webb, 1998]:
  - $\text{Eval}_i(x, y)$: ask agent i for $v_i([x, y])$
  - $\text{Cut}_i(x, \alpha)$: ask agent for a point $y \in [0,1]$ such that $v_i([x, y]) = \alpha$

Existence of Fair Allocations

Direct Revelation:

- Partition $[0, 1]$ to $\{X_1, \ldots, X_m\}$ s. t. each f_i is uniform on each X_j.
- Allocate each X_j uniformly to all the n agents.
- Envy-free and Proportional: $\forall i, j: v_i(A_j) = \frac{1}{n} v_i([0,1])$

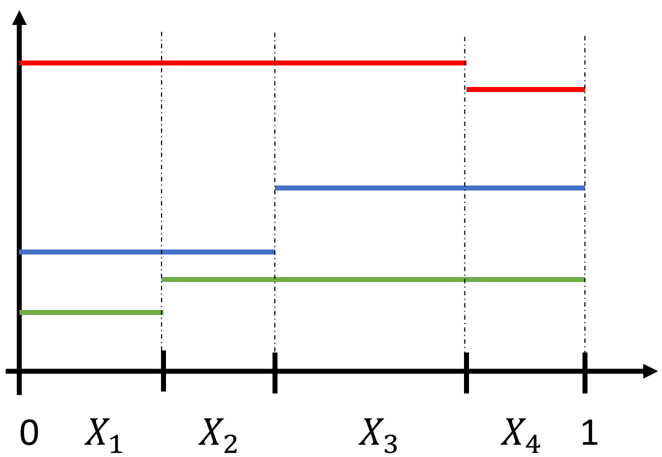

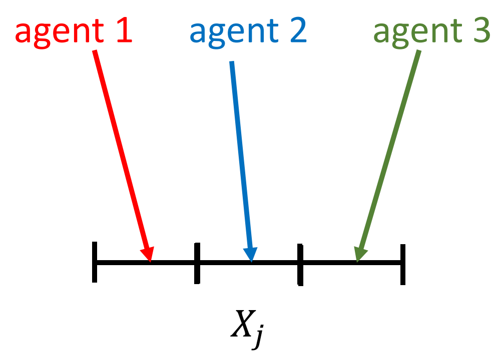

Existence of Fair Allocations

Robertson-Webb Query Model:

- Proportional allocations always exist!
  - Moving-knife procedure [Dubins & Spanier, 1961]: $\Theta(n^2)$ queries
  - Even & Paz (1984): $\Theta(n \log n)$ queries
  - Lower-bound: $\Omega(n \log n)$ queries [Edmonds & Pruhs, 2006]
- Envy-Free allocations always exist!
  - [Aziz & Machenzie, 2017]: $n^{n^{n^{n^{n^{n}}}}}$ queries!
  - Lower bound: $\Omega(n^2)$ queries [Procaccia, 2009]

Mechanism

- Mechanism: $\mathcal{M}: (f_1, \ldots, f_n) \rightarrow (A_1, \ldots, A_n)$
  - deterministic

- Agents report their valuations on the cake to the mechanism
- Mechanism decides an allocation upon receiving the reports

Truthful Mechanism

- Agents may lie and misreport their valuations if beneficial! → Game Theory
- Truthful mechanism: mechanism under which reporting truthfully is a dominant strategy.
- "Whatever other agents do, truth-telling is always the best strategy for me"
- For any $f_1, \ldots, f_n$ and any f_1',
- Let $\mathcal{M}(f_1, f_2, \ldots, f_n) = (A_1, \ldots, A_n)$ and $\mathcal{M}(f_1', f_2, \ldots, f_n) = (A_1', \ldots, A_n')$
- Truthful $\implies$ $v_1(A_1) \geq v_1(A_1')$

Objective

- A fair (envy-free or proportional) and truthful mechanism $\mathcal{M}$?

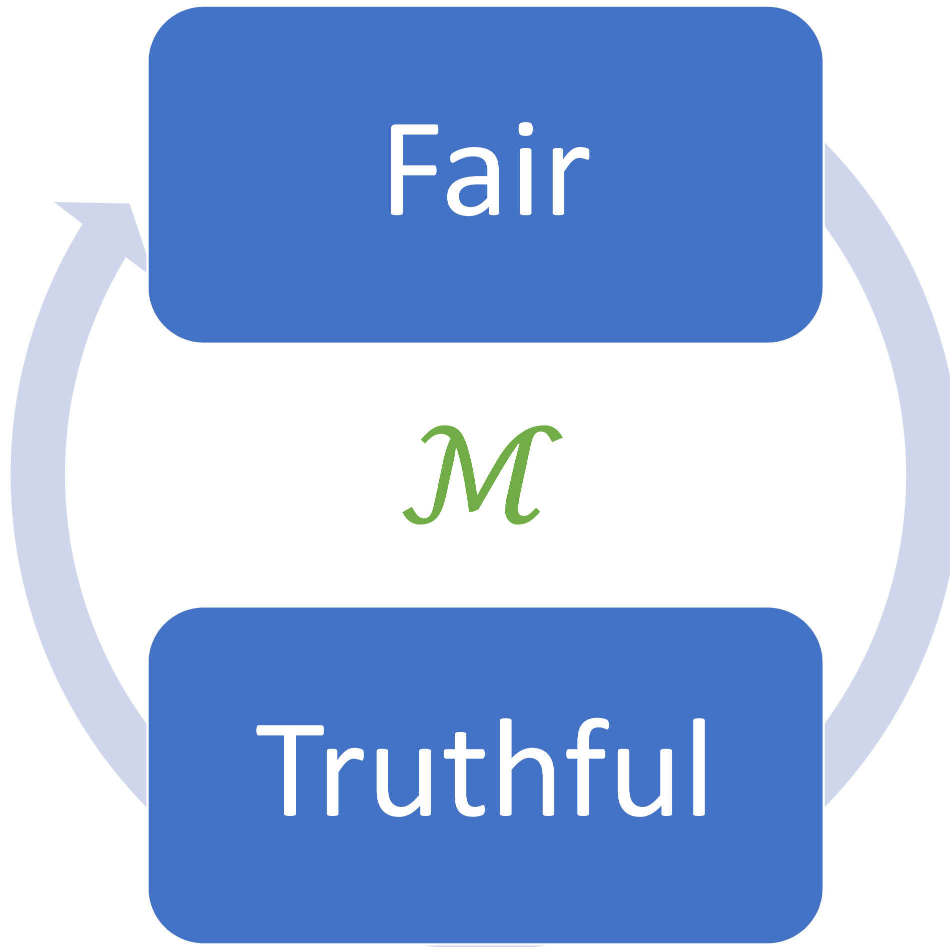

Mechanism: *I-cut-you-choose*

- "I-cut-you-choose": a mechanism for two agents.
- An agent A cuts the cake into two parts with equal value (w.r.t. his/her value density function)
- The other agent B chooses one part, and A gets the other part.

Example

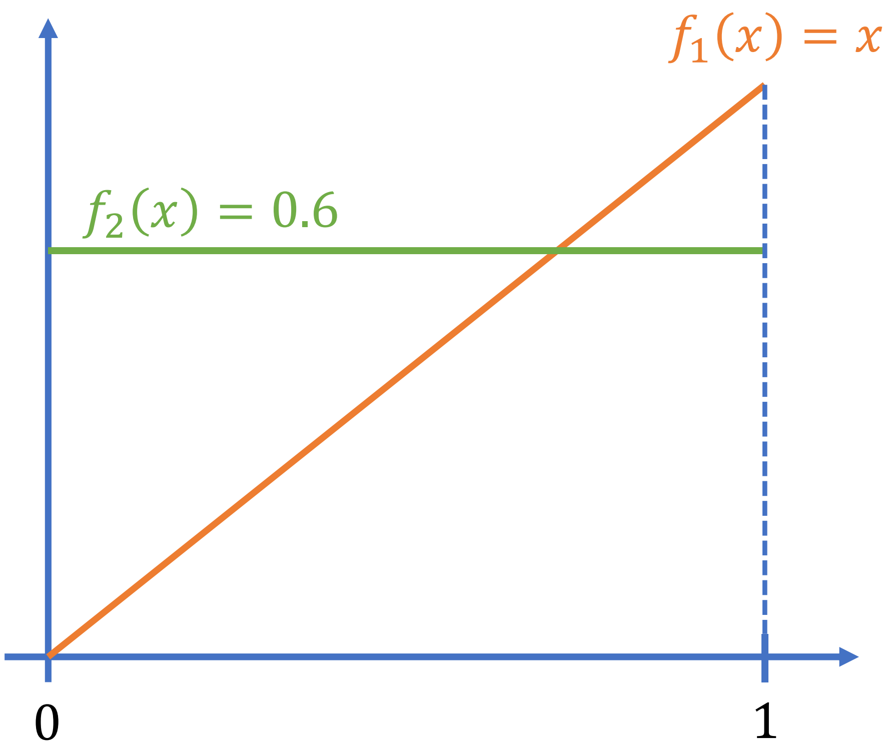

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the “cutter”.

Example

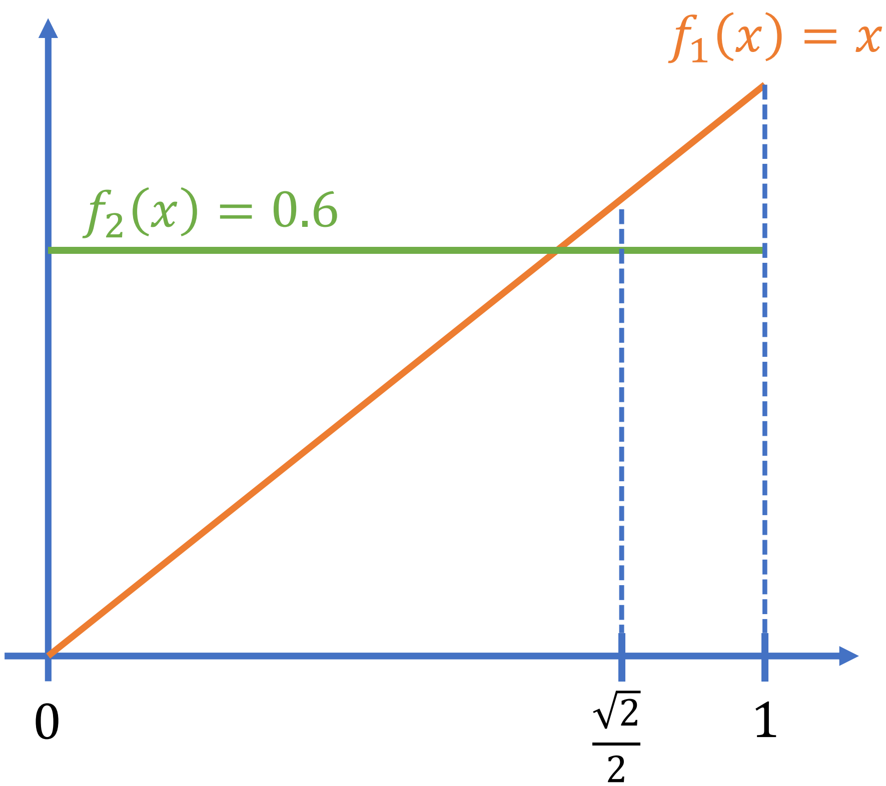

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the "cutter".
- Agent 1 will cut at $x = \frac{\sqrt{2}}{2}$.
- Both $[0, x)$ and $[x, 1]$ worth $\frac{1}{4}$.

Example

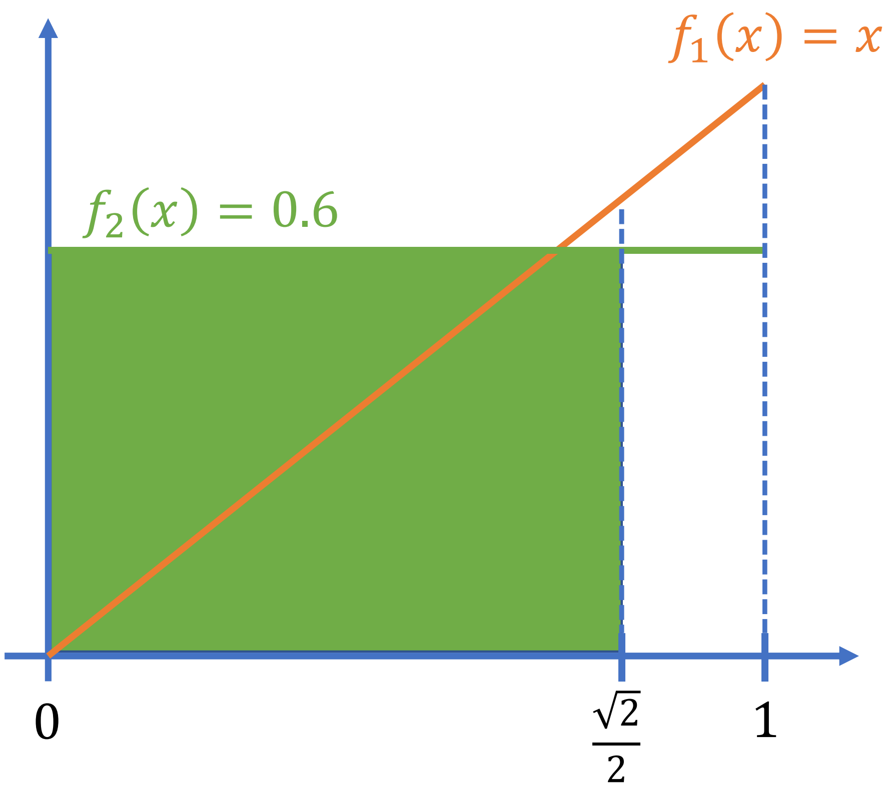

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the “cutter”.
- Agent 1 will cut at $x = \frac{\sqrt{2}}{2}$.
- Both $[0, x)$ and $[x, 1]$ worth $\frac{1}{4}$.
- Agent 2 will choose $[0, x)$, which has a larger value.

Example

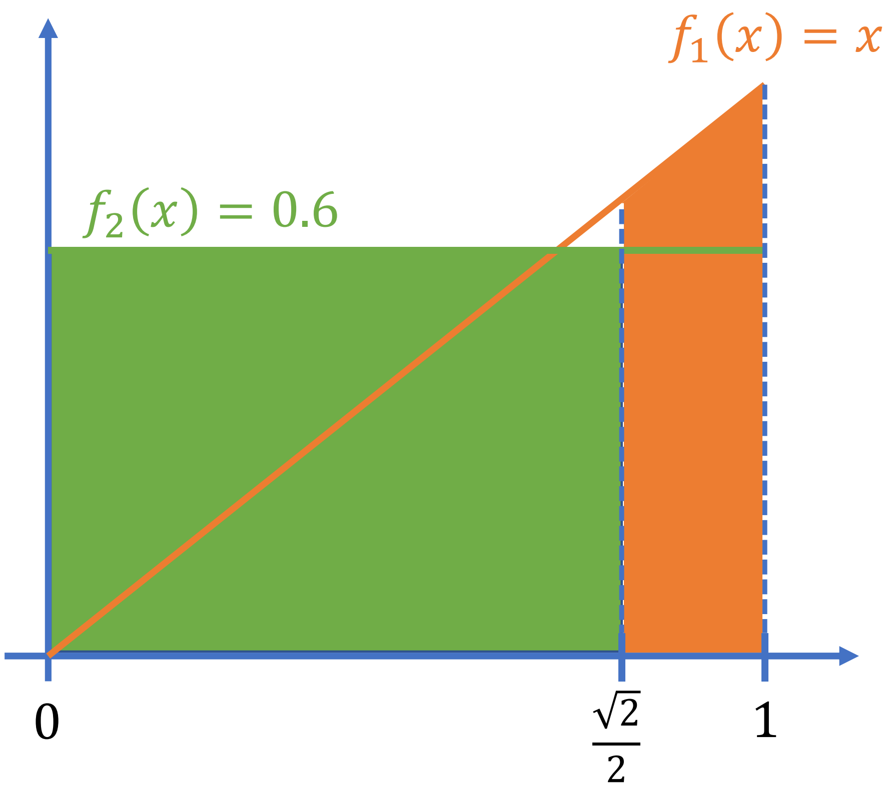

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the "cutter".
- Agent 1 will cut at $x = \frac{\sqrt{2}}{2}$.
- Both $[0, x)$ and $[x, 1]$ worth $\frac{1}{4}$.
- Agent 2 will choose $[0, x)$, which has a larger value.
- $[x, 1]$ is left to agent 1.

Mechanism: *I-cut-you-choose*

- "I-cut-you-choose": a mechanism for two agents.
- An agent A cuts the cake into two parts with equal value (w.r.t. his/her value density function)
- The other agent B chooses one part, and A gets the other part.
- "I-cut-you-choose" is envy-free.
- Envy-free $\Longrightarrow$ proportional
- Is it truthful?

Example

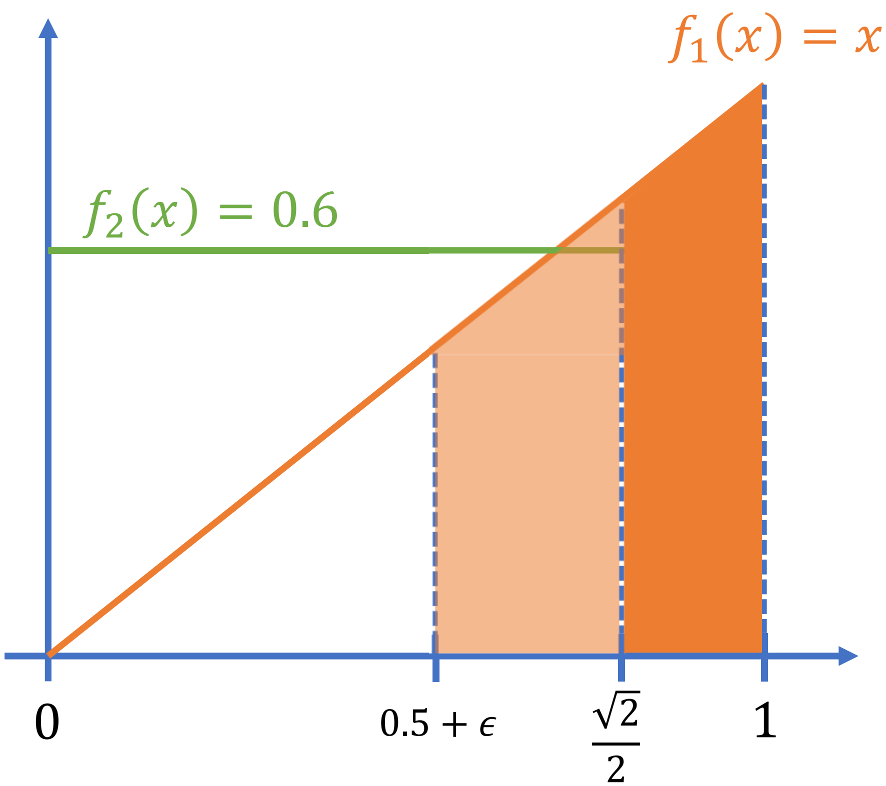

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the “cutter”.
- Agent 1 would like to cut at $x = 0.5 + \epsilon$, so that (s)he will receive more value.

Example

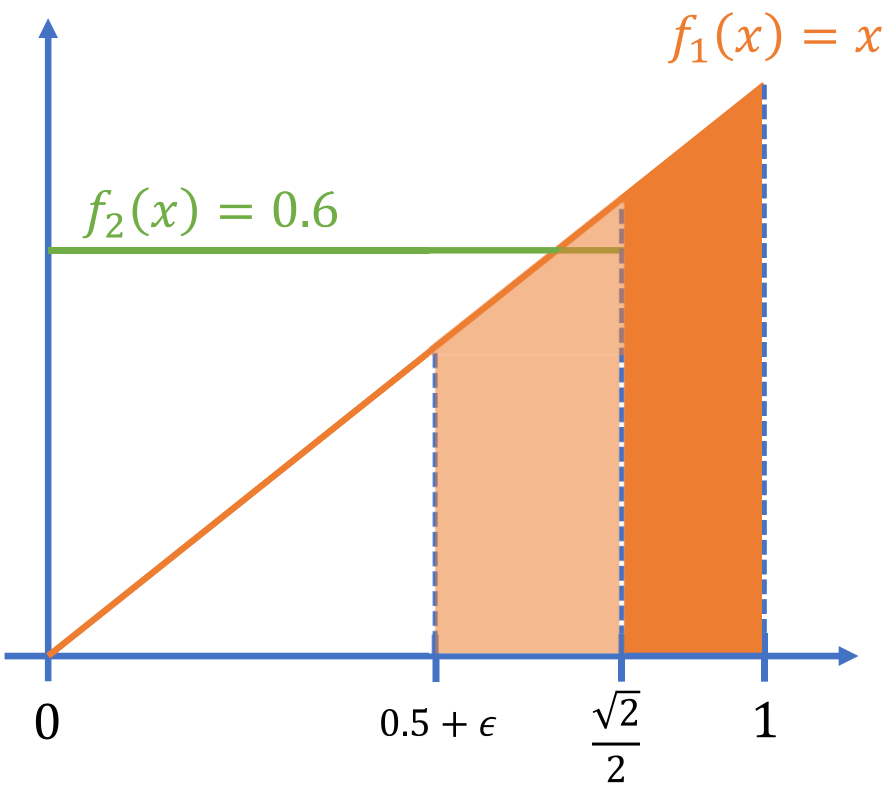

- Value density functions:
  - Agent 1: $f_1(x) = x$
  - Agent 2: $f_2(x) = 0.6$
- Suppose agent 1 is the "cutter".
- Agent 1 would like to cut at $x = 0.5 + \epsilon$, so that (s)he will receive more value.
- Agent 1 can report a function f_1' such that $0.5 + \epsilon$ is the half-half point.

Mechanism: *I-cut-you-choose*

- "I-cut-you-choose": a mechanism for two agents.
- An agent A cuts the cake into two parts with equal value (w.r.t. his/her value density function)
- The other agent B chooses one part, and A gets the other part.
- "I-cut-you-choose" is envy-free.
- Envy-free $\Longrightarrow$ proportional
- Is it truthful? **NO!!!**

Truthful Mechanism

- Agents may lie and misreport their valuations if beneficial! → Game Theory
- Truthful mechanism: mechanism under which reporting truthfully is a dominant strategy.

- **Question:** Does there exist a truthful fair (envy-free or proportional) cake cutting mechanism?

Two Game-Theoretical Settings

- Direct Revelation:
  - assuming each f_i is piecewise constant
  - $\mathcal{M}: \{f_1, \dots, f_n\} \to \{A_1, \dots, A_n\}$
  - a one-round game
- Robertson-Webb query model:
  - extensive-form game
  - agents can act adaptively based on previous queries
  - more room for manipulation

Impossibility Results for Robertson-Webb Query Model

- [Kurokawa et al., 2013] There is no truthful and envy-free mechanism with a bounded number of queries.
- [Brânzei and Miltersen, 2015] For any truthful mechanism, there exists an agent who receives a zero value.

Direct Revelation (Piecewise Uniform VDFs)

- Piecewise uniform f_i:
  - $\{X_1, \ldots, X_m\}$: partition of $[0, 1]$, f_i is either 0 or 1 on each X_t
- [Chen et al., 2010] a truthful envy-free mechanism, with free-disposal assumption
- [Maya & Nisan, 2012] characterization of truthful envy-free mechanisms, uniqueness of Chen et al.'s mechanism.
- [Li et al., 2015] Chen et al.'s mechanism also works under settings with externalities.
- [Bei et al., 2020] a truthful envy-free mechanism, without free-disposal assumption.

Direct Revelation (Piecewise Constant VDFs)

- [Aziz & Ye, 2014] No truthful mechanism satisfies both proportionality and Pareto-optimal.
- [Menon & Larson, 2017] No truthful mechanism satisfies even approximately-proportional, with the connected pieces constraint.
- [Bei et al., 2017] No truthful, proportional mechanism exists under <u>any</u> of the following settings:
  - non-wasteful
  - position-oblivious
  - sequential reporting VDFs

Main Result

- Question: Does there exist a truthful, proportional cake-cutting mechanism?
- This work: NO!
- There does not exist a truthful, proportional mechanism.
- This impossibility result even holds under the following settings:
  - There are only two agents.
  - Discarding some parts of the cake is allowed.
  - Agents' valuations on the cake are always positive.

Proof Ideas

- Suppose such a mechanism $\mathcal{M}$ exists.
- Construct a sequence of cake-cutting instances.
- Show that truthfulness and proportionality cannot be both satisfied on all the constructed instances.

Instance F^1

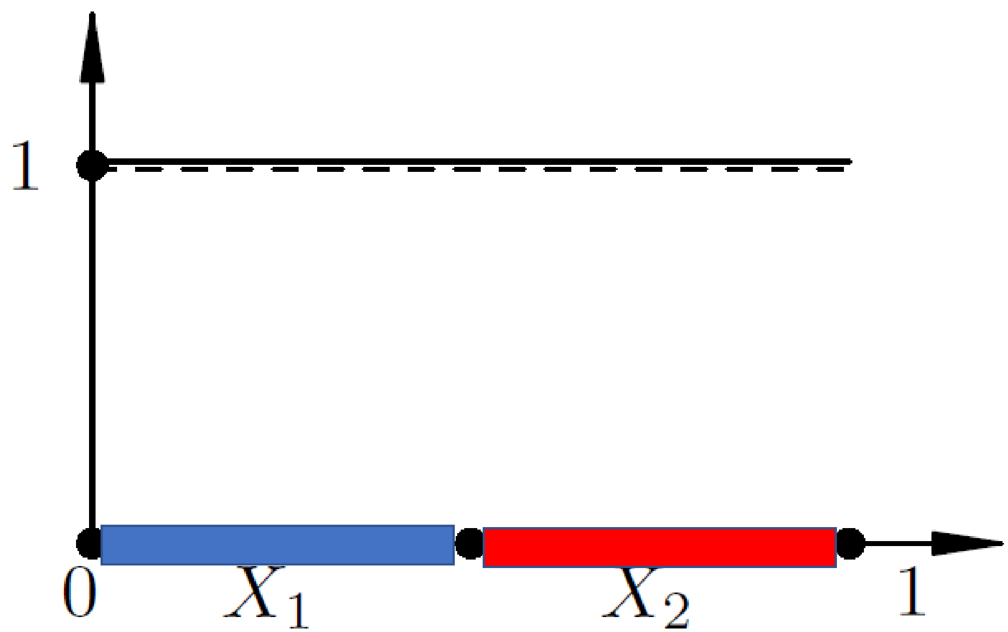

- $f_1(x) = 1$
- $f_2(x) = 1$

- $\mathcal{M}(F^1) = (X_1, X_2)$ with
- $|X_1| = |X_2| = 0.5$

Instance F^2

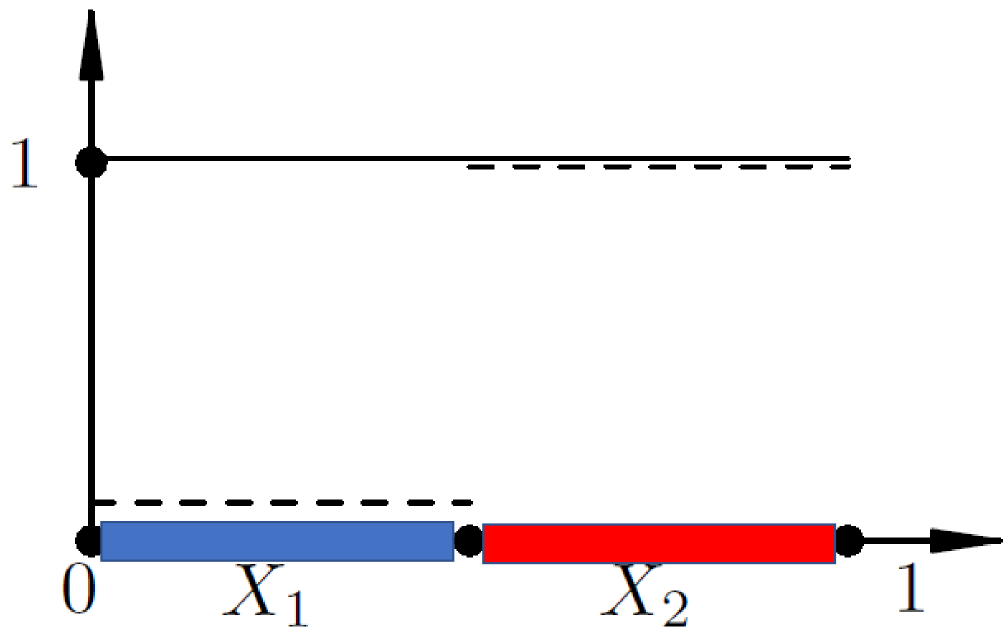

- $f_1(x) = 1$
- $f_2(x) = \begin{cases} \epsilon, & x \in X_1 \\ 1, & x \in X_2 \end{cases}$

- $\mathcal{M}(F^2) = (X_1, X_2)$ with
- $|X_1| = |X_2| = 0.5$
- the same allocation with $\mathcal{M}(F^1)$

Instance F^3

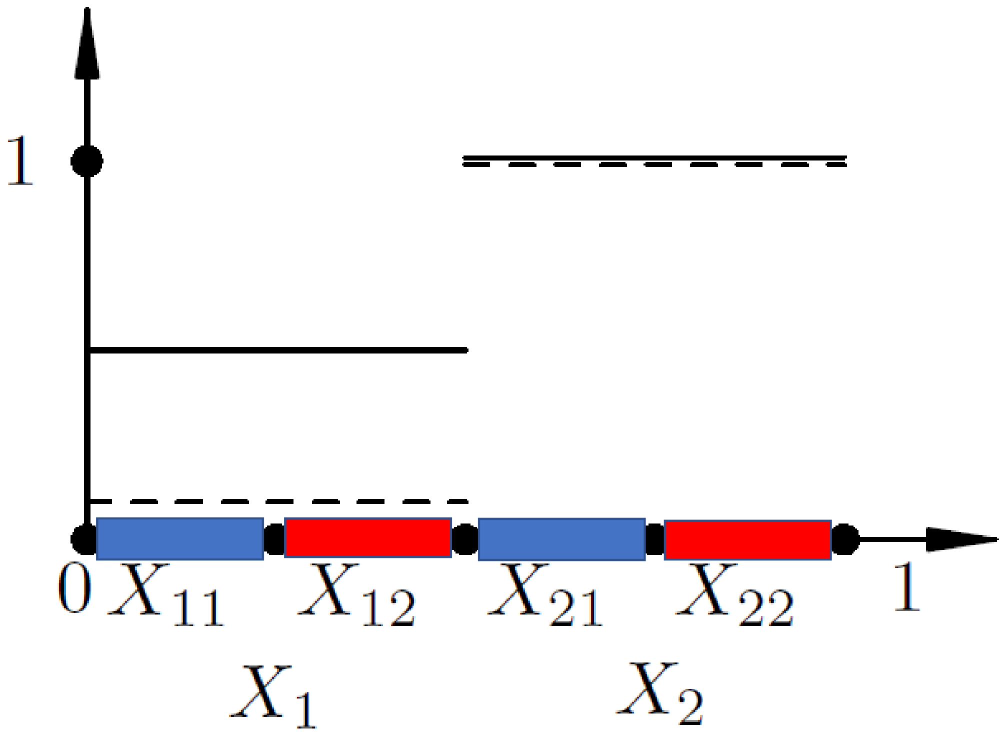

- Agent 1 cannot receive a total length of more than 0.5
  - O.W., in F^2, he will misreport to F^3
- Agent 1 cannot receive less than half of X_2
  - O.W., the total length he receives is more than 0.5, to guarantee proportionality
- Agent 1 cannot receive more than half of X_2
  - O.W., cannot guarantee proportionality
- $\mathcal{M}\left(F^3\right) = (X_{11} \cup X_{21}, X_{12} \cup X_{22})$, with $|X_{11}| = |X_{12}| = |X_{21}| = |X_{22}| = 0.25$

Remaining part of the proof: web of instances

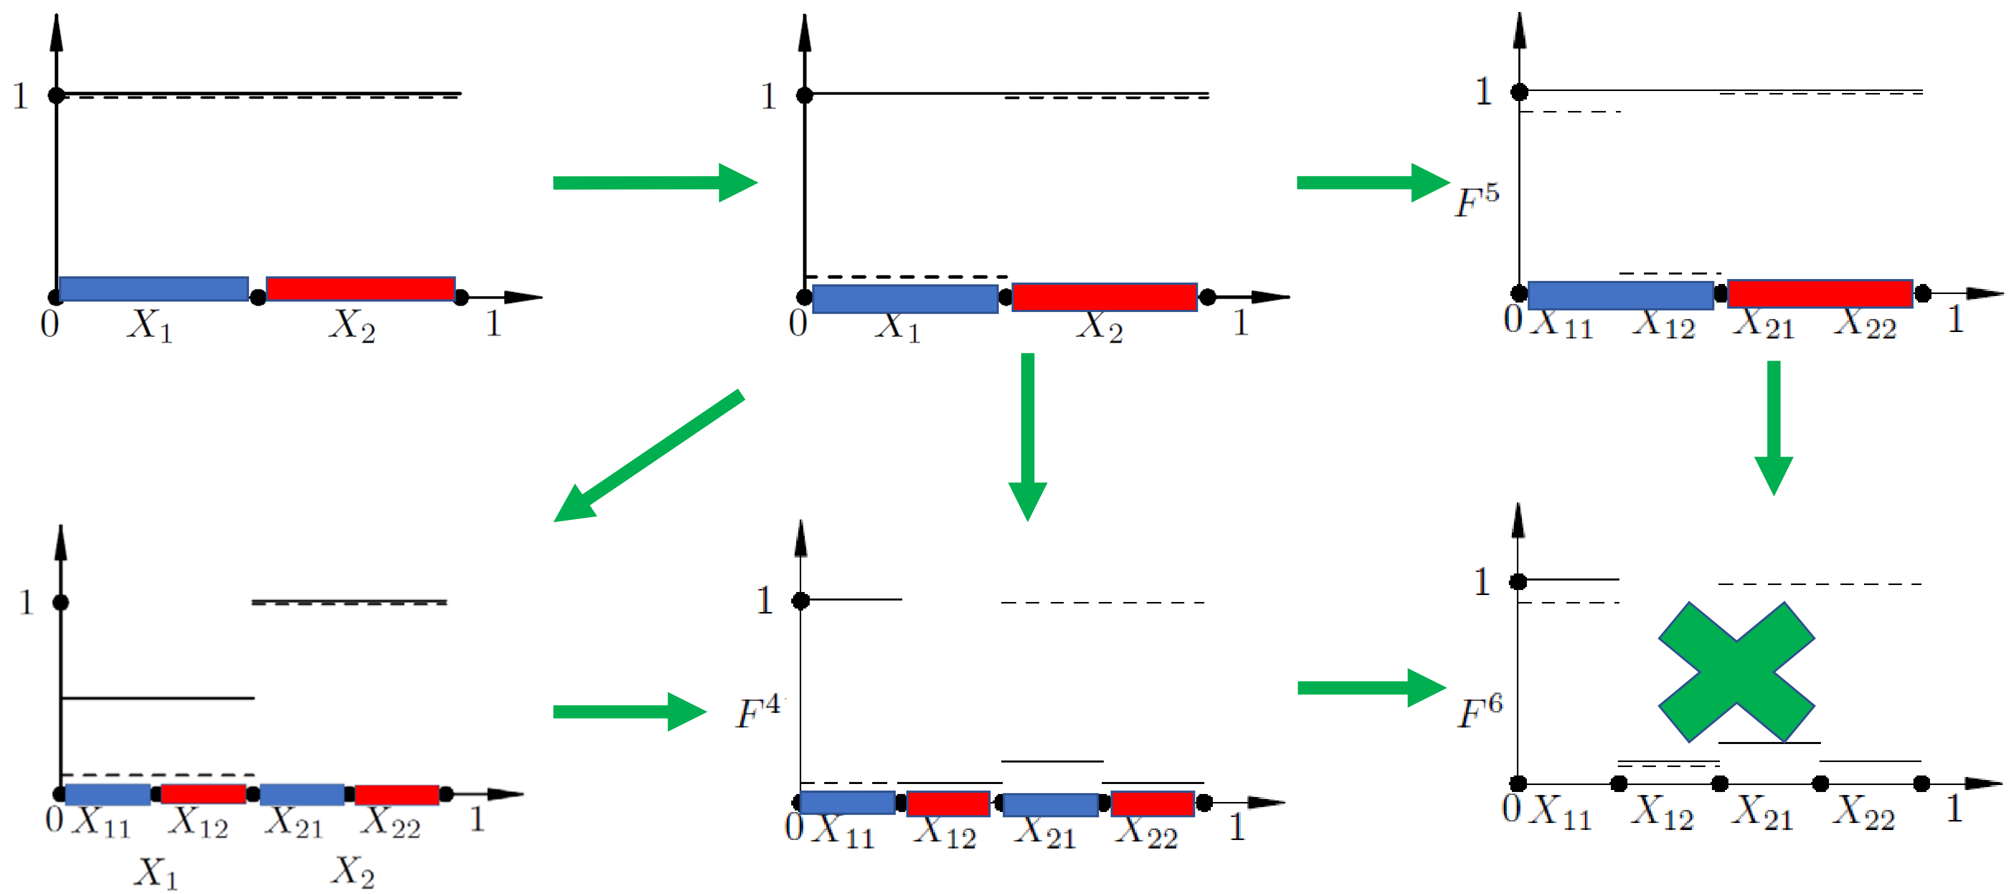

Circumventing the Impossibility Result

- Relaxing truthfulness
- Relaxing fairness (proportionality)

Relaxing truthfulness

- Dominant-strategy truthful: “Whatever other agents do, truth-telling is always the best strategy for me.”
- Any weaker truthful notions?
- Truth-telling as a Nash Equilibrium: “If other agents report valuations truthfully, truth-telling is the best strategy for me.”
- Unfortunately, in the context of cake cutting, the two notions are equivalent.
- We need even weaker truthful notions…

Motivation

- We have seen "I-cut-you-choose" is not (dominant-strategy) truthful.
- However, it still can achieve "some kind of truthfulness".
- It is dominant-strategy truthful for the chooser.
- Dominant-strategy truthfulness fails for the cutter.
- However, manipulating the cutting point is risky if the cutter does not know the value density function of the chooser!
- There is always a possibility that the cutter ends up getting a part with a less-than-proportional value.

Risk-Averse Truthfulness

$\mathcal{M}$ is risk-averse truthful if

- either misreporting f_i is non-beneficial
- or there is a chance $\mathcal{M}$ allocate A_i to agent i with $v_i(A_i) < \frac{1}{n} v_i([0,1])$

- I-cut-you-choose is risk-averse truthful.
- Can we extend I-cut-you-choose to general number of agents?

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- Let each agent i cut the cake into n intervals each of which has value $\frac{1}{n} v_i([0, 1])$.
- Each agent i introduces $n - 1$ cut points $x_1^{(i)}, x_2^{(i)}, \dots, x_{n-1}^{(i)}$.
- $S = \{1, \dots, n\}$ be the set of agents who have not been allocated
- For $k = 1, \dots, n - 1$:
  - Find the left-most cut point from the k-th set of cut points $\left\{x_k^{(i)}\right\}_{i \in S}$
  - Let i_k be the agent whose k-th cut point is at left-most among agents in S
  - Allocate the part to the left of the cut point to i_k, and remove i_k from S
- Allocate the remaining part of the cake to the one remaining agent

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- Example with three agents: blue, orange, and green.

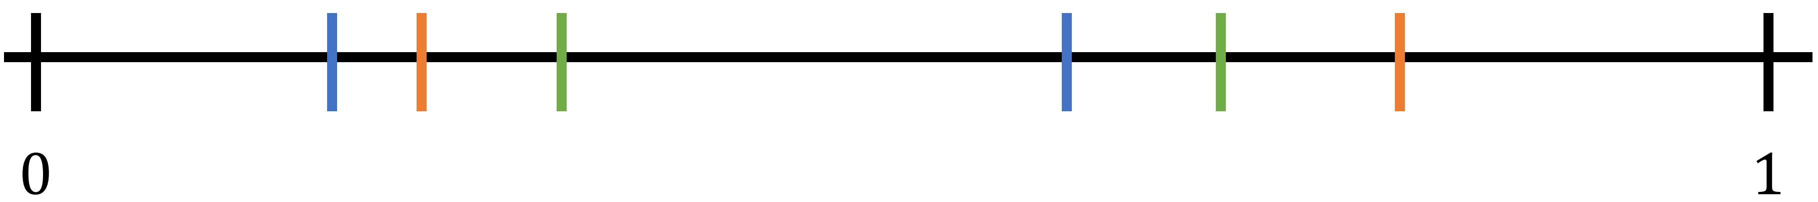

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- For the first set of cut points, the blue one is at the left-most.

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- For the first set of cut points, the blue one is at the left-most.
- Allocation for the blue agent.

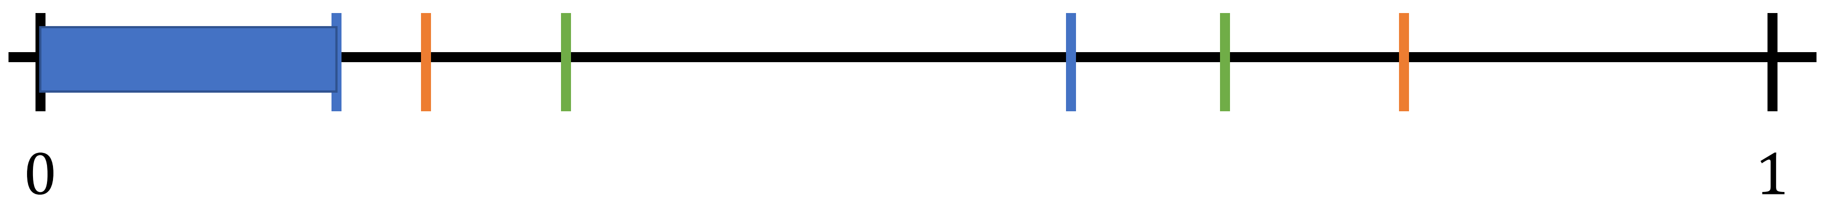

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- For the first set of cut points, the blue one is at the left-most.
- Allocation for the blue agent.
- Remove the blue agent.

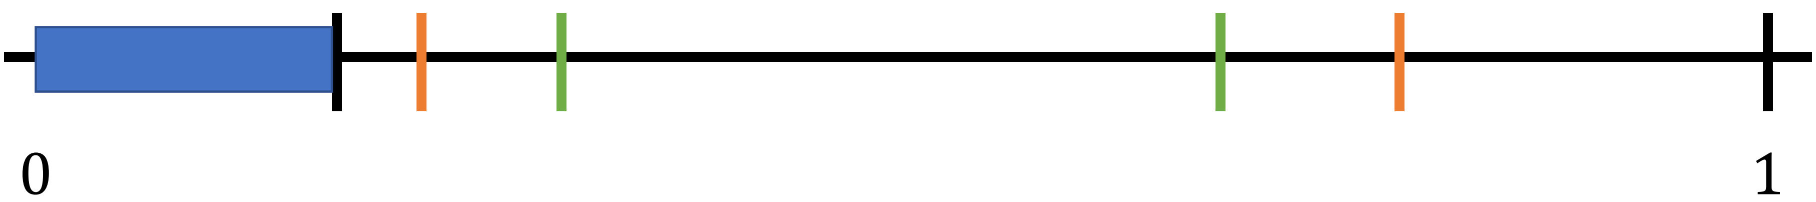

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- For the second set of cut points, the green one is at the left-most.

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- For the second set of cut points, the green one is at the left-most.
- Allocation for the green agent.
- Remove the green agent.

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- The remaining part is allocated to the remaining agent---the orange agent.

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- $\mathcal{M}_{PROP}$ is proportional!
- The first agent receives value exactly $\frac{1}{n} v_i([0,1])$.
- Each remaining agent receives an interval that contains two cut points.
- Thus, each remaining agent receives value at least $\frac{1}{n} v_i([0,1])$.

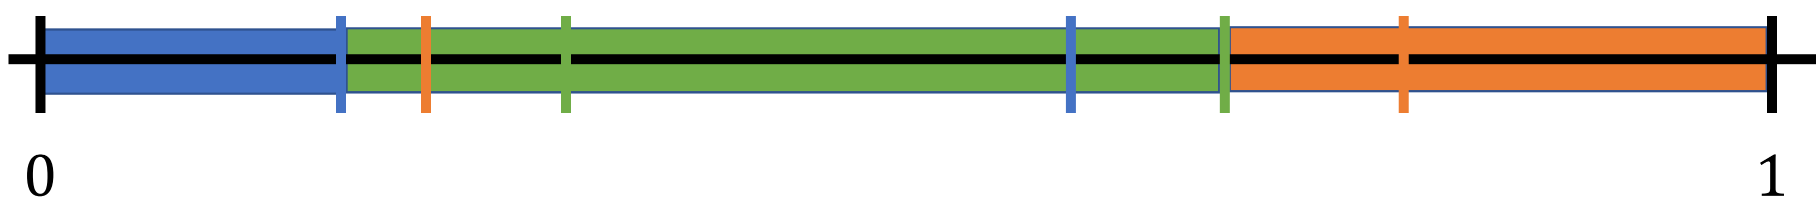

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- $\mathcal{M}_{PROP}$ is risk-averse truthful!
- If reporting f_i' instead of f_i does not change the position of the cut points, the misreporting is non-beneficial.

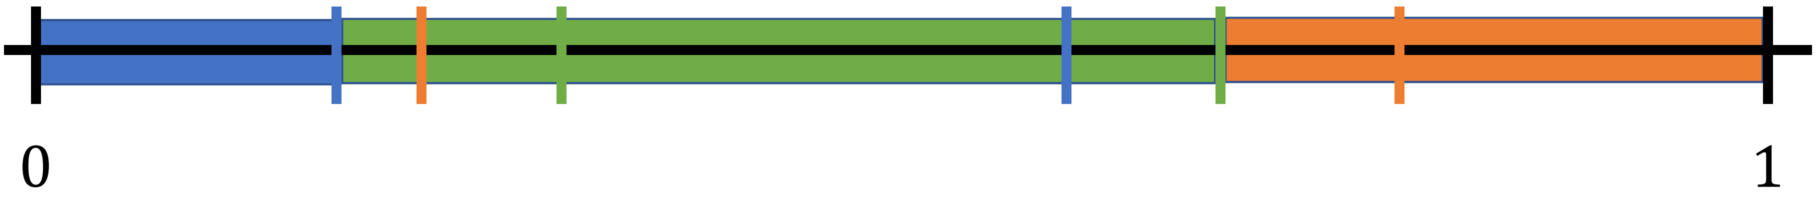

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- $\mathcal{M}_{PROP}$ is risk-averse truthful!
- If reporting f_i' instead of f_i does not change the position of the cut points, the misreporting is non-beneficial.
- Otherwise, one of the n intervals has value less than $\frac{1}{n}v_i([0,1])$

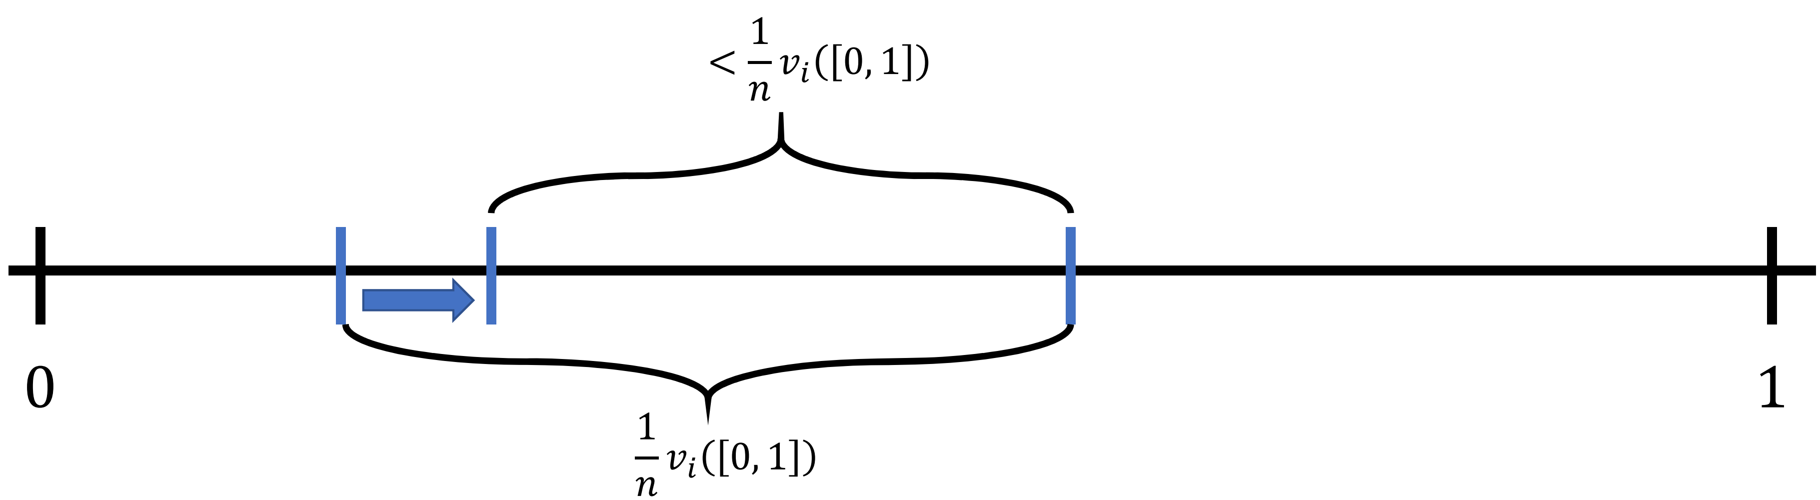

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- $\mathcal{M}_{PROP}$ is risk-averse truthful!
- If reporting f_i' instead of f_i does not change the position of the cut points, the misreporting is non-beneficial.
- Otherwise, one of the n intervals has value less than $\frac{1}{n}v_i([0,1])$.
- There is a chance that this agent will receive this interval.

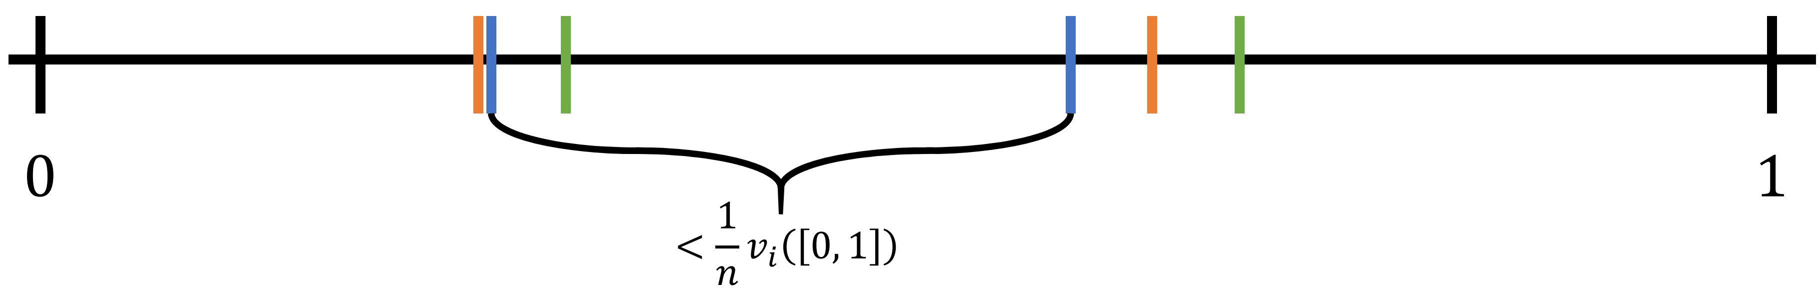

A proportional and risk-averse truthful mechanism $\mathcal{M}_{PROP}$

- $\mathcal{M}_{PROP}$ is risk-averse truthful!
- If reporting f_i' instead of f_i does not change the position of the cut points, the misreporting is non-beneficial.
- Otherwise, one of the n intervals has value less than $\frac{1}{n}v_i([0,1])$.
- There is a chance that this agent will receive this interval.

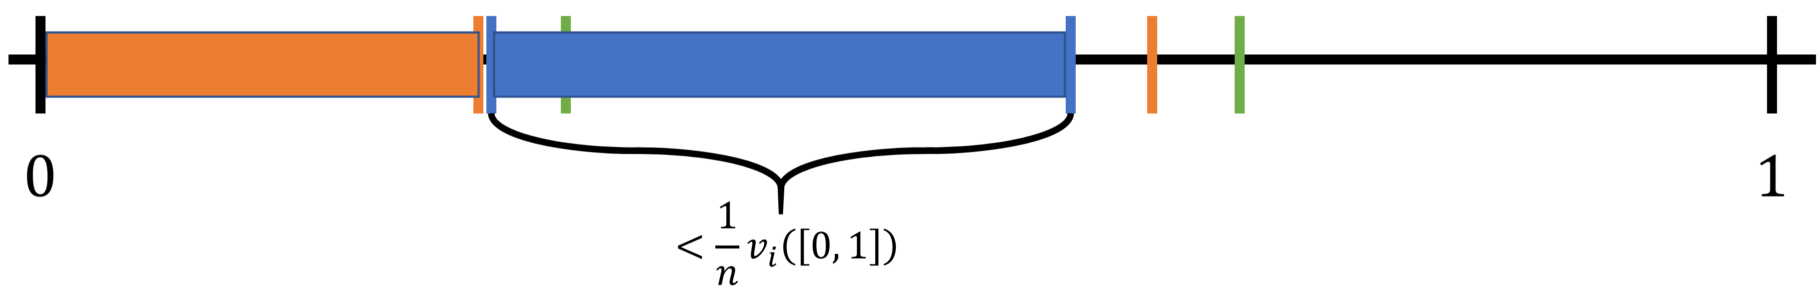

Envy-Free + risk-averse truthful?

- How about an envy-free and risk-averse truthful mechanism?

Envy-Free + risk-averse truthful?

- Let's revisit the envy-free mechanism introduced earlier.
- This mechanism is envy-free, but is it risk-averse truthful?
- **NO**! Not if the order of the n agents are fixed on each X_j.

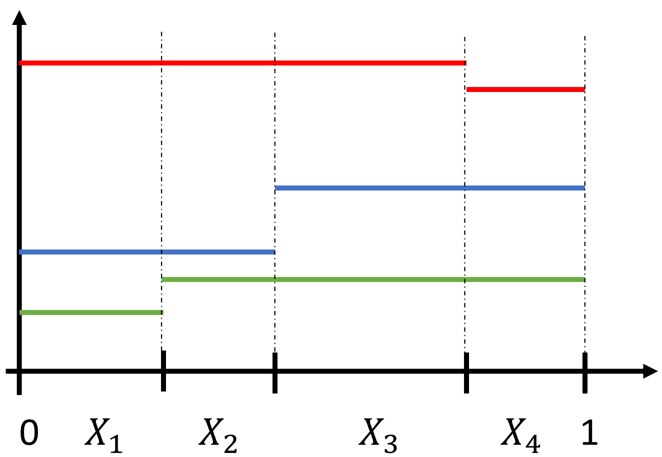

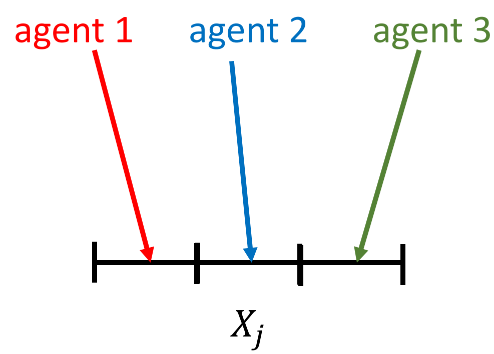

Counterexample

- Suppose $f_1(x) = \begin{cases} 2, & x \in [0, 0.5) \\ 1, & x \in (0.5, 1] \end{cases}$.

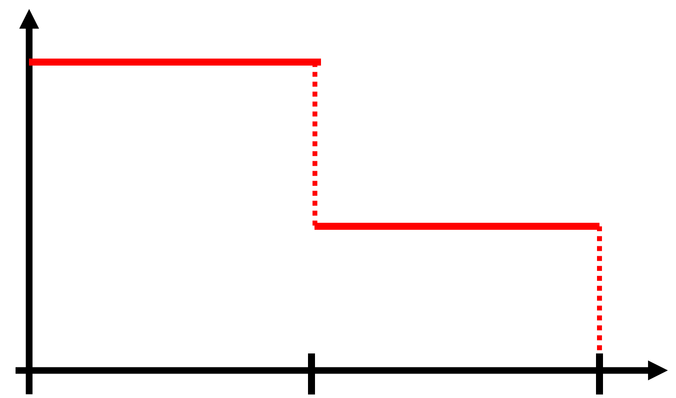

Counterexample

- Suppose $f_1(x) = \begin{cases} 2, & x \in [0, 0.5) \\ 1, & x \in (0.5, 1] \end{cases}$.
- Reporting $f_1'(x) = 2$ is always no worse and sometimes better!

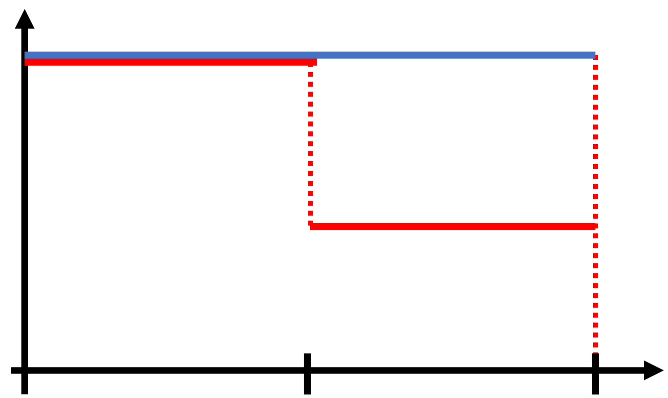

Counterexample

- Suppose $f_1(x) = \begin{cases} 2, & x \in [0, 0.5) \\ 1, & x \in (0.5, 1] \end{cases}$.
- Reporting $f_1'(x) = 2$ is always no worse and sometimes better!
- The break-point 0.5 is eliminated with f_1'

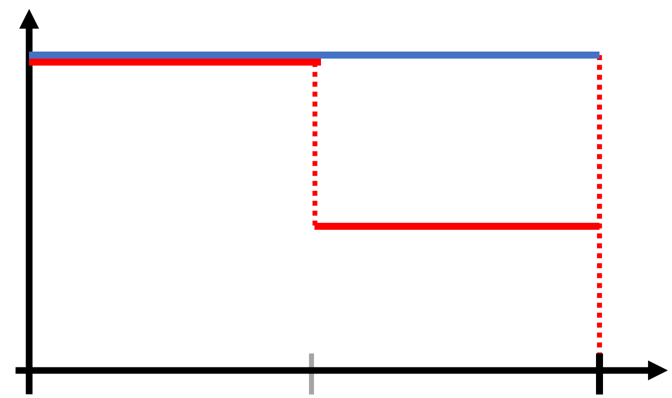

Counterexample

- Suppose $f_1(x) = \begin{cases} 2, & x \in [0, 0.5) \\ 1, & x \in (0.5, 1] \end{cases}$.
- Reporting $f_1'(x) = 2$ is always no worse and sometimes better!
- The break-point 0.5 is eliminated with f_1'
- No matter how the cake is partitioned $\{X_j\}$

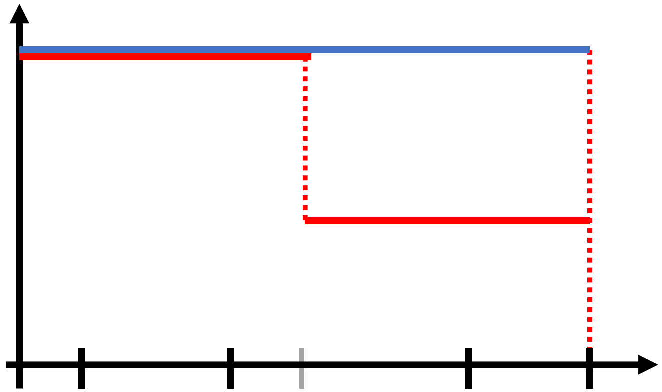

Counterexample

- Suppose $f_1(x) = \begin{cases} 2, & x \in [0, 0.5) \\ 1, & x \in (0.5, 1] \end{cases}$.
- Reporting $f_1'(x) = 2$ is always no worse and sometimes better!
- The break-point 0.5 is eliminated with f_1'
- No matter how the cake is partitioned $\{X_j\}$
- Eliminating a break point is always beneficial, since agent 1 is the left-most agent!

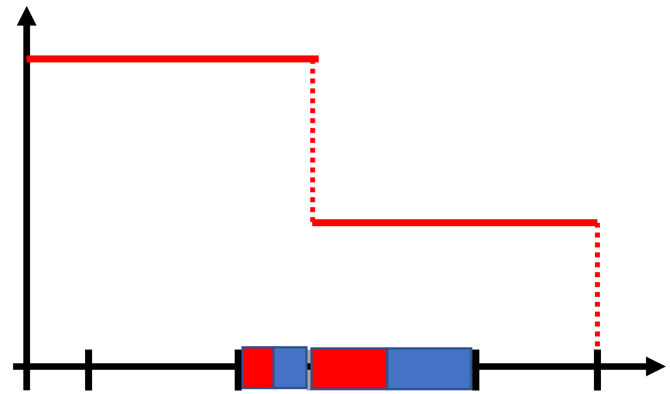

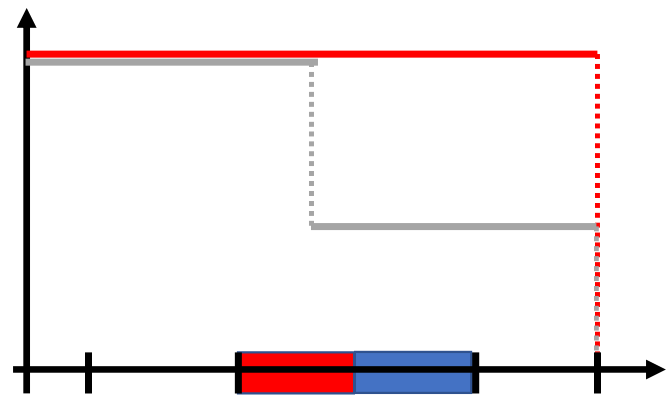

Envy-Free + risk-averse truthful?

- Let's revisit the envy-free mechanism introduced earlier.
- This mechanism is envy-free, but is it risk-averse truthful?
- **NO**! Not if the order of the n agents are fixed on each X_j.
- The fixed order is problematic!

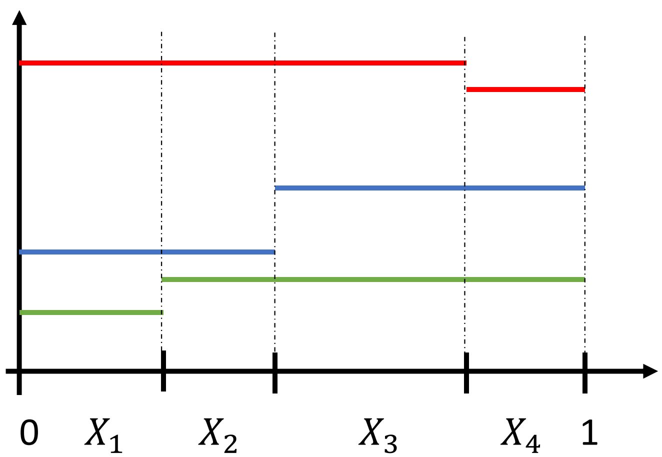

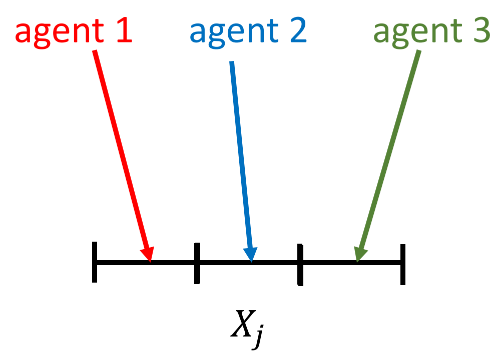

An envy-free and risk-averse truthful mechanism $\mathcal{M}_{EF}$

- Similar to the previous mechanism
- Different order in different X_j
- E.g., set agent $i = j \pmod{n}$ is the left-most agent on X_j
- Risk-averse truthful now: agents cannot foresee the order on X_j

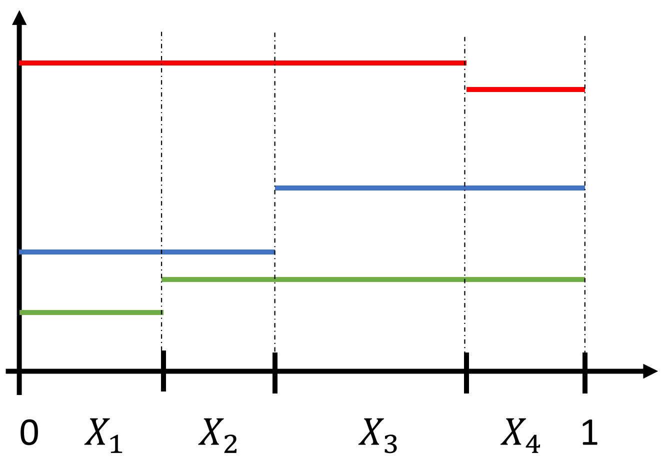

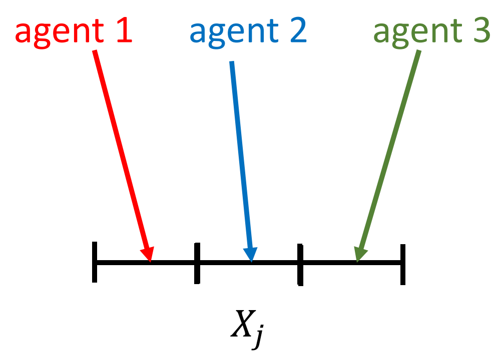

Comparison

$\mathcal{M}_{EF}$:

- Envy-free
- Proportional
- Risk-averse truthful

$\mathcal{M}_{PROP}$:

- ~~Envy-free~~
- Proportional
- Risk-averse truthful
- Output connected-pieces allocations

Relaxing Fairness

- α-approximate proportionality:

$$\forall i\colon\ v_i(A_i) \geq \alpha \cdot \frac{1}{n} v_i([0,1])$$

Result

- The main impossibility result extends to the setting with (exact) proportionality relaxed to 0.974031-approximate proportionality.

Future Work

- Does there exist $\alpha > 0$ such that there exists a truthful, α-approximately proportional mechanism?
- Does there exist a truthful mechanism that is non-oblivious to each agent's VDF?
- What a truthful mechanism can do?

Thank You!

A Randomized Envy-Free Mechanism [Mossel & Tamuz, 2010]

- Partition $[0,1]$ to $\{X_1, \ldots, X_m\}$ s. t. each f_i is uniform on each X_j.
- Allocate each X_j randomly and uniformly to all the n agents.

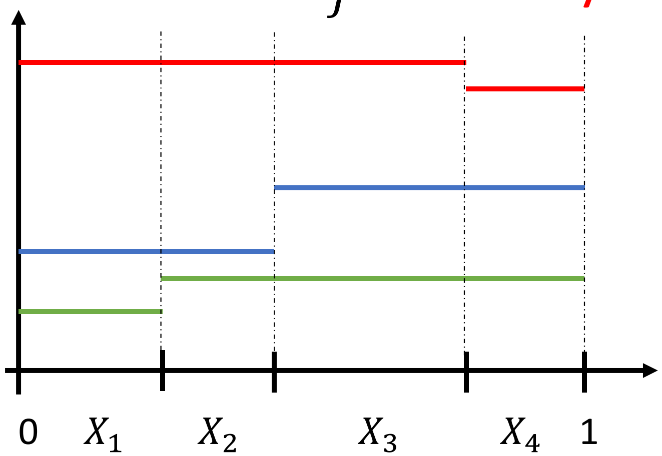

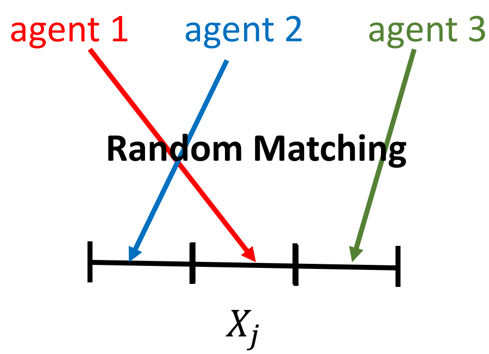

- Universal envy-free and truthful in expectation.