TRUTHFUL AGGREGATION OF BUDGET PROPOSALS

Rupert Freeman

UVA Darden School of Business

Based on joint work with Mark de Berg, Javier Cembrano, David M. Pennock, Dominik Peters, Ulrike Schmidt-Kraepelin, Markus Utke, Jennifer Wortman Vaughan

Budget Aggregation

- A common budget or resource needs to be (fully) divided among m projects
- Each project can receive an arbitrary percentage
- n voters report their most preferred budget

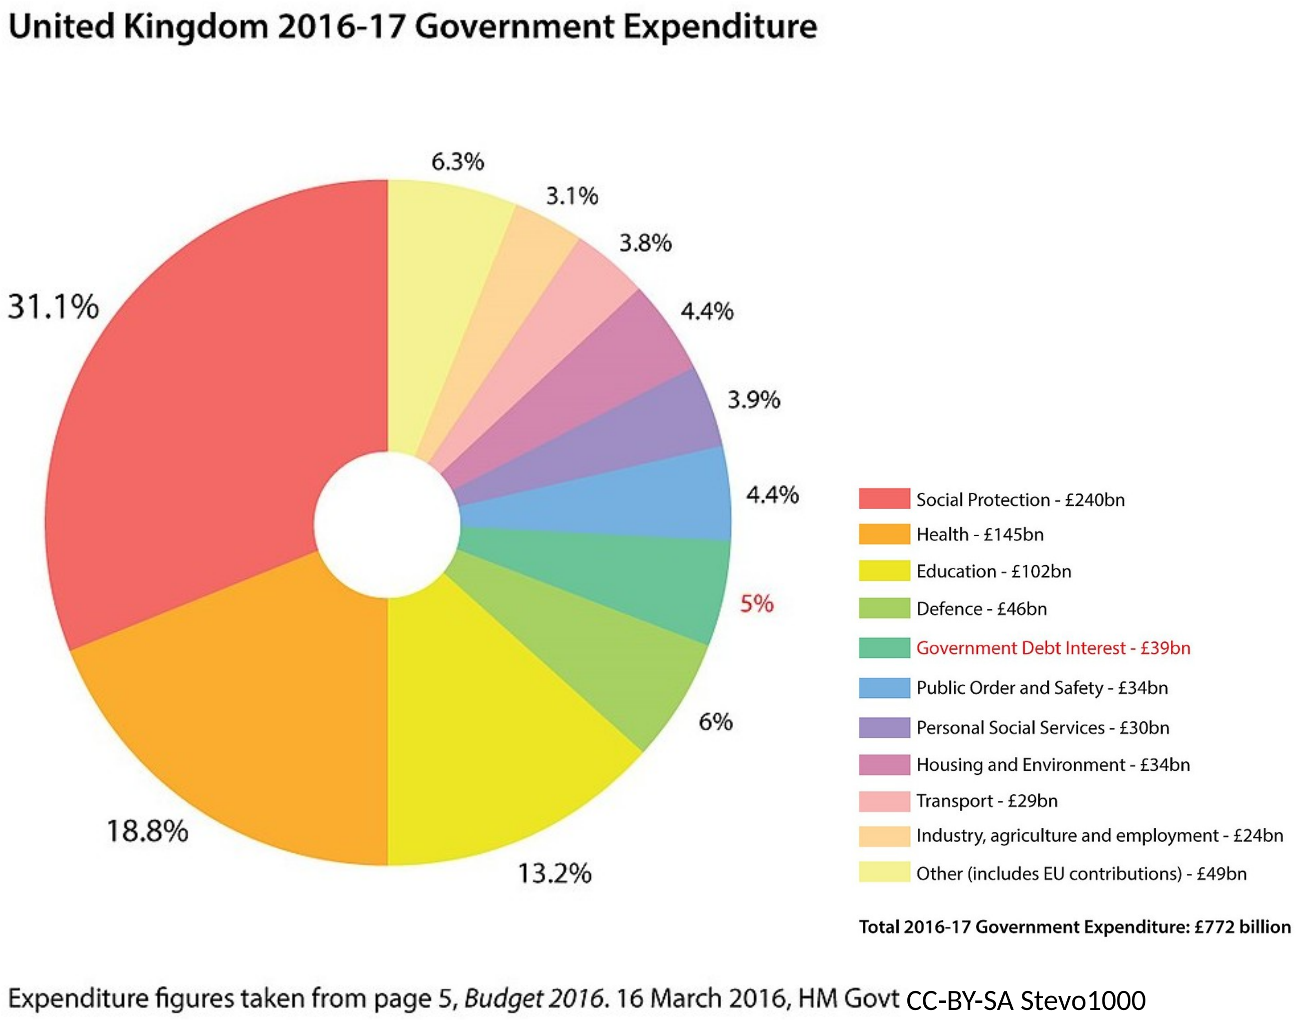

Example

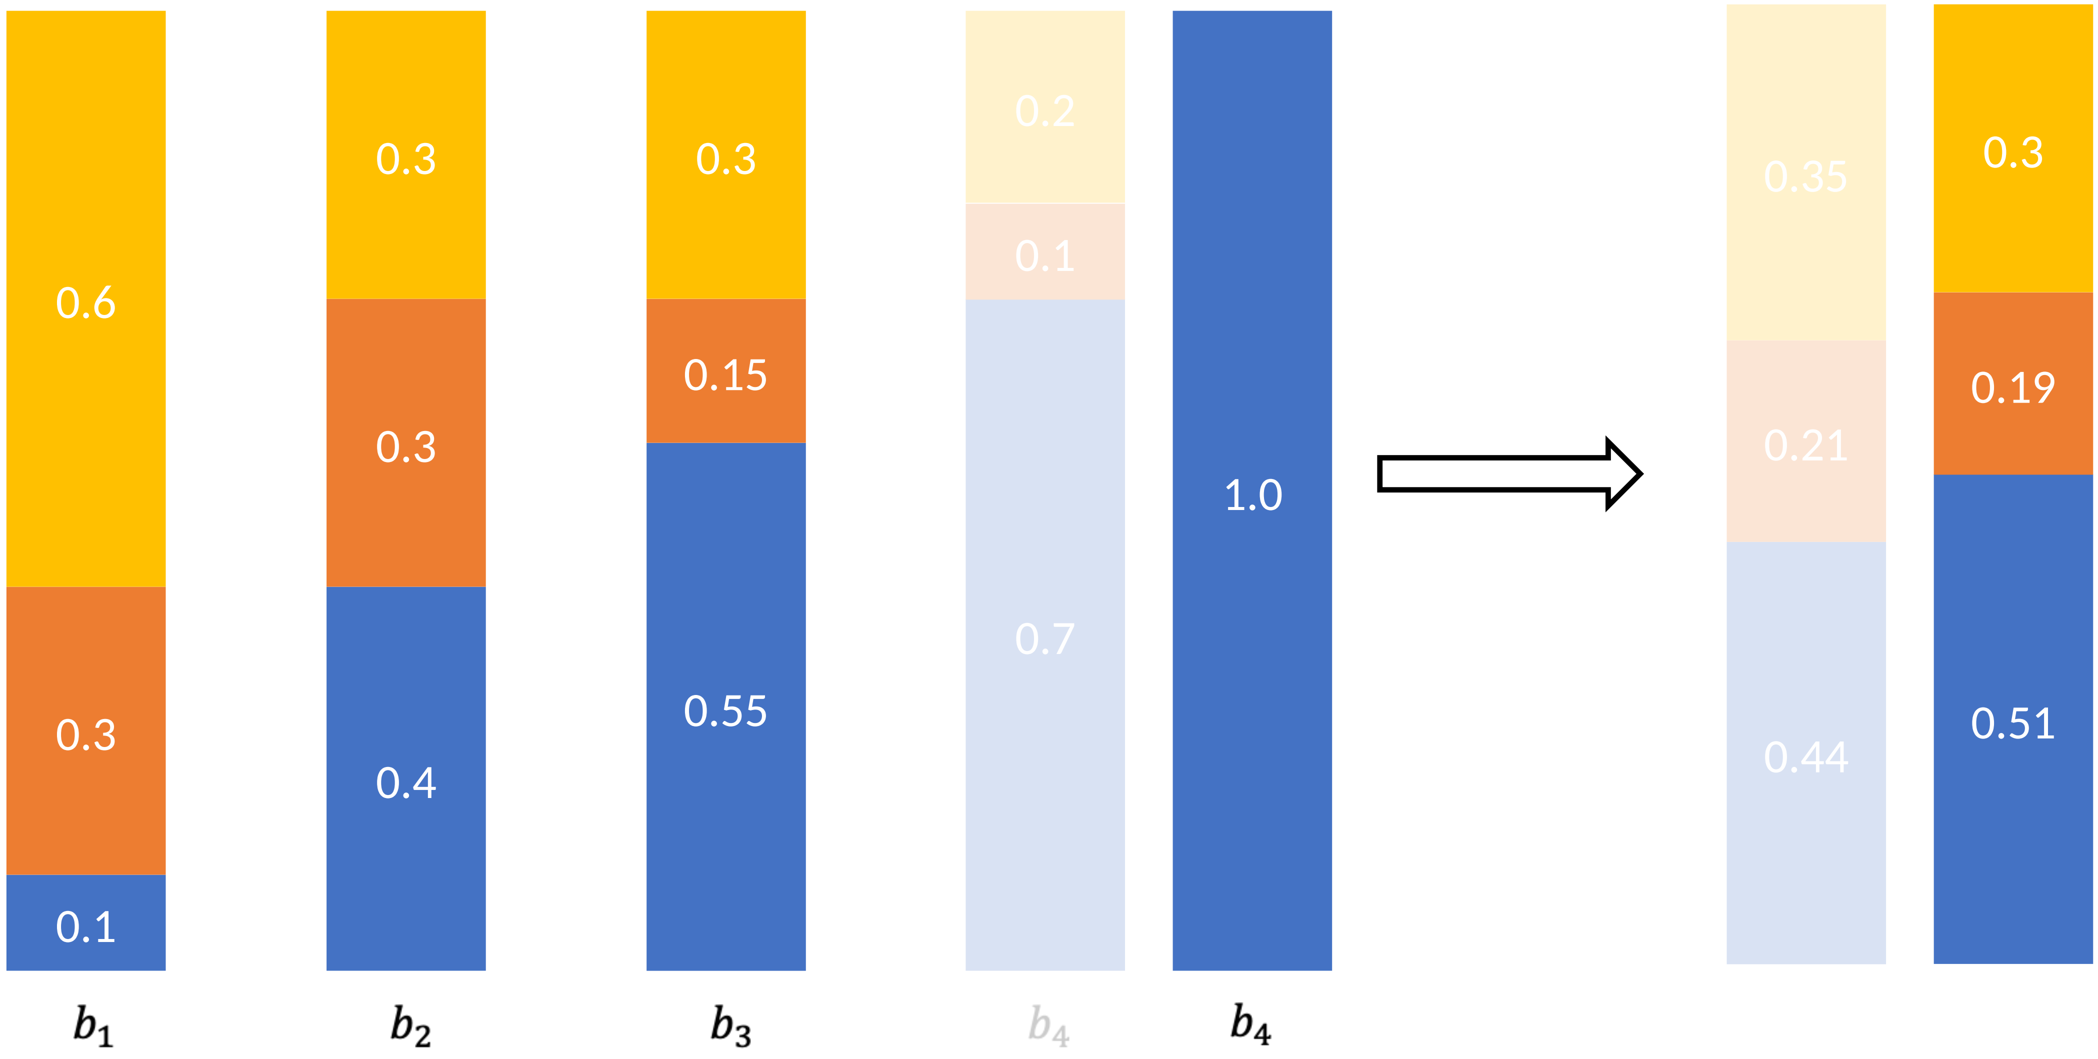

0.6
0.3
0.1
0.3
0.3
0.4
0.3
0.15
0.55
0.2
0.1
0.7
1.0
0.35
0.21
0.44
0.3
0.19
0.51
b_1
b_2
b_3
b_4
b_4

Model

- Voters prefer distributions that are closer to their ideal
  - we use ℓ_1 metric in this work: $d(p,q) = \sum_{j \in [m]} |p_j - q_j|$

[Nehring and Puppe 2019]

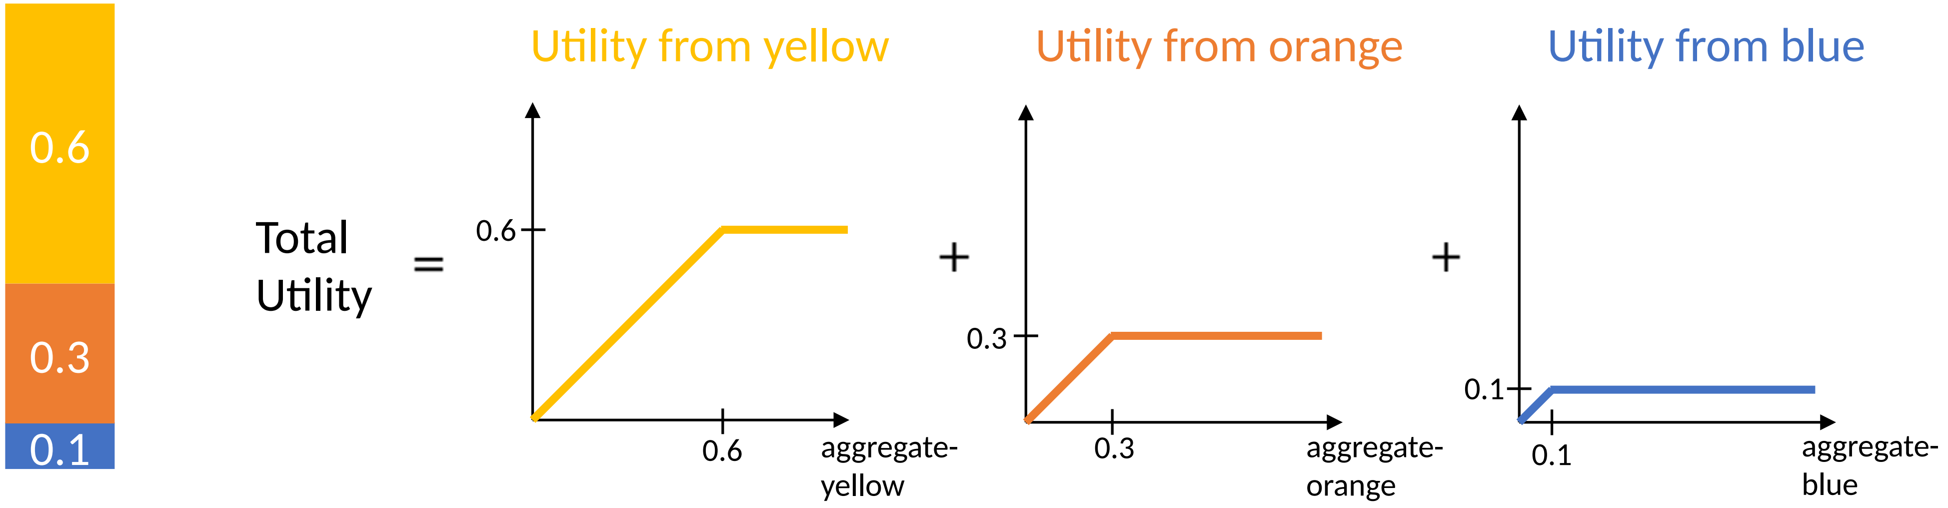

- Strategyproof mechanism: No voter can benefit from reporting a distribution that is not her ideal

A Strategyproof Mechanism

- Util: Select the distribution that maximizes social welfare (minimizes sum of distances to votes)
- [Lindner et al. 2008, Goel et al. 2019]: Util is strategyproof.
- But not (single-minded) proportional

If all voters are ***single-minded*** then the resource should be distributed proportionally to each project‘s support.

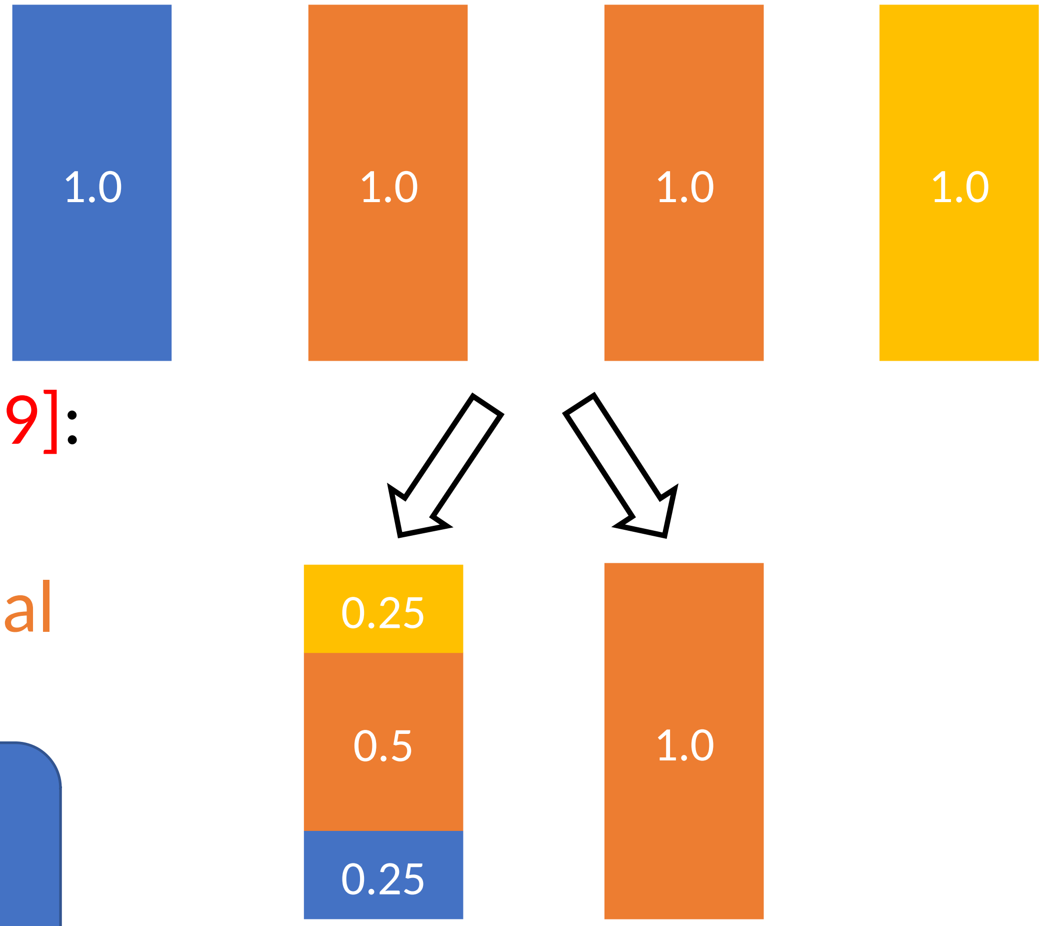

Outline

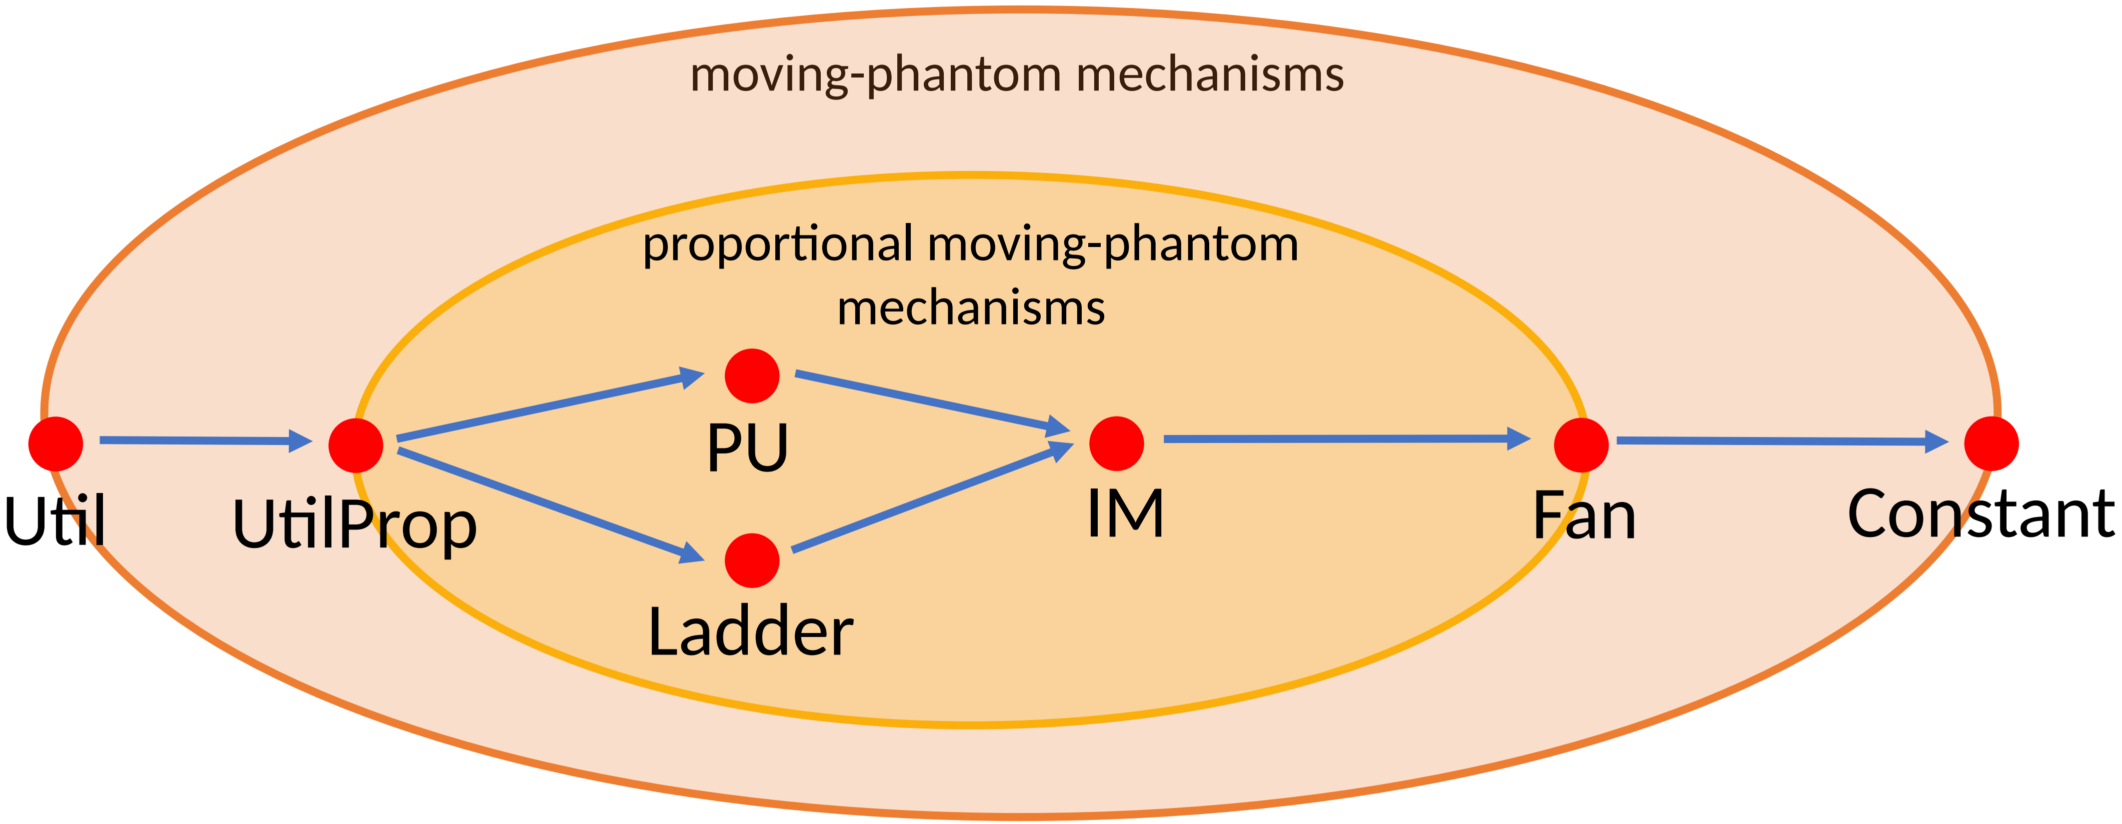

moving-phantom mechanisms
proportional moving-phantom mechanisms
Util
UtilProp
PU
Ladder
IM
Fan
Constant

Two Projects

Two Projects

- Suppose 3 voters report:
  - (0.3, 0.7)
  - (0.1, 0.9)
  - (0.5, 0.5)

- Median is normalized and strategyproof

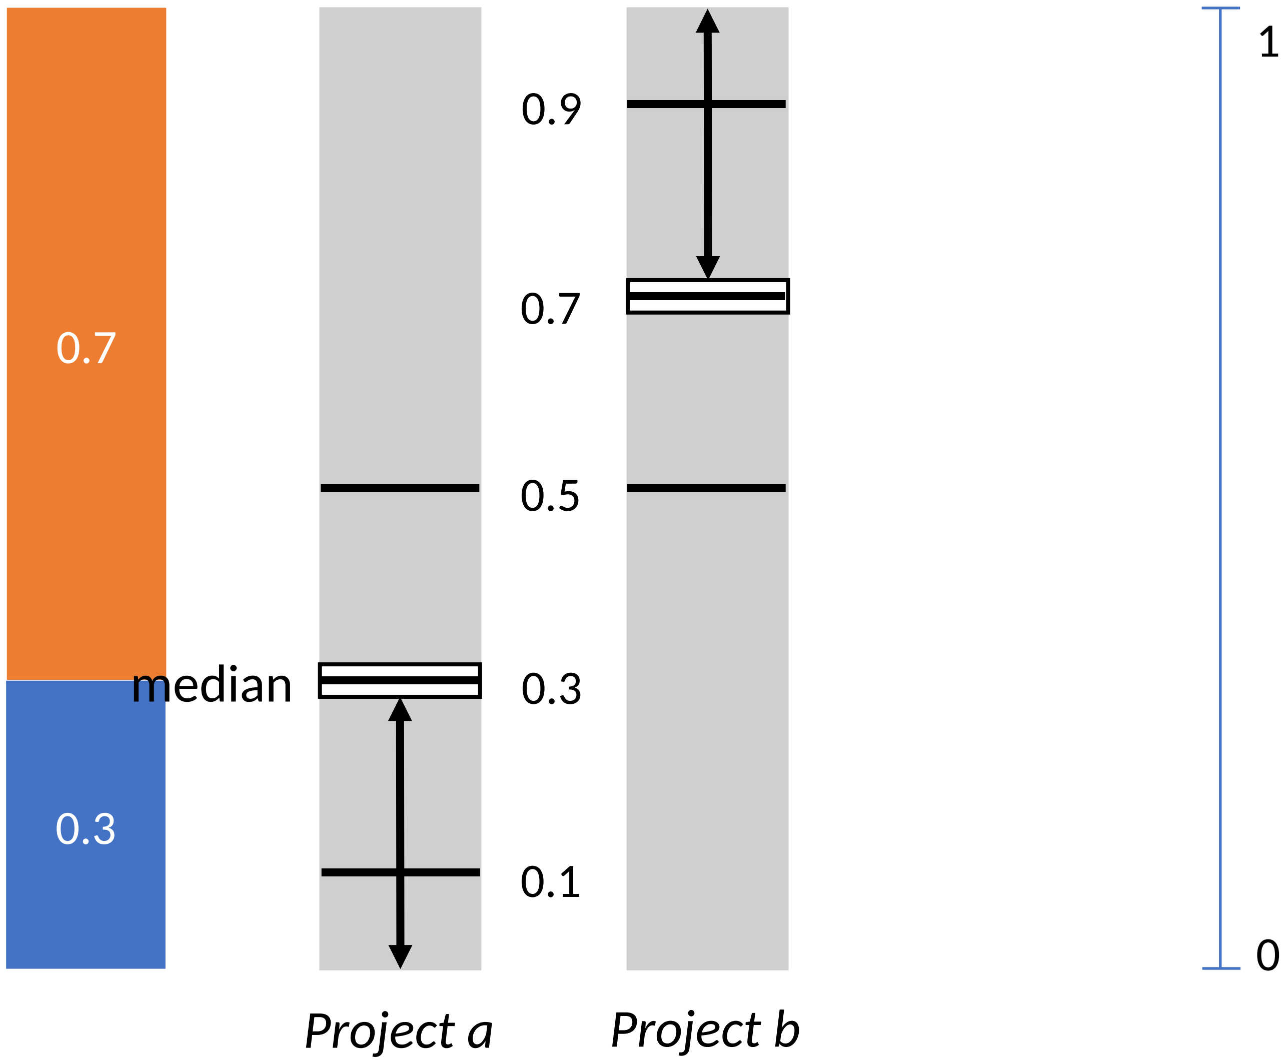

Two Projects

- Suppose 3 voters report:
  - (0.3, 0.7)
  - (0.1, 0.9)
  - (0.5, 0.5)
- Median is normalized and strategyproof
- What else can we do?

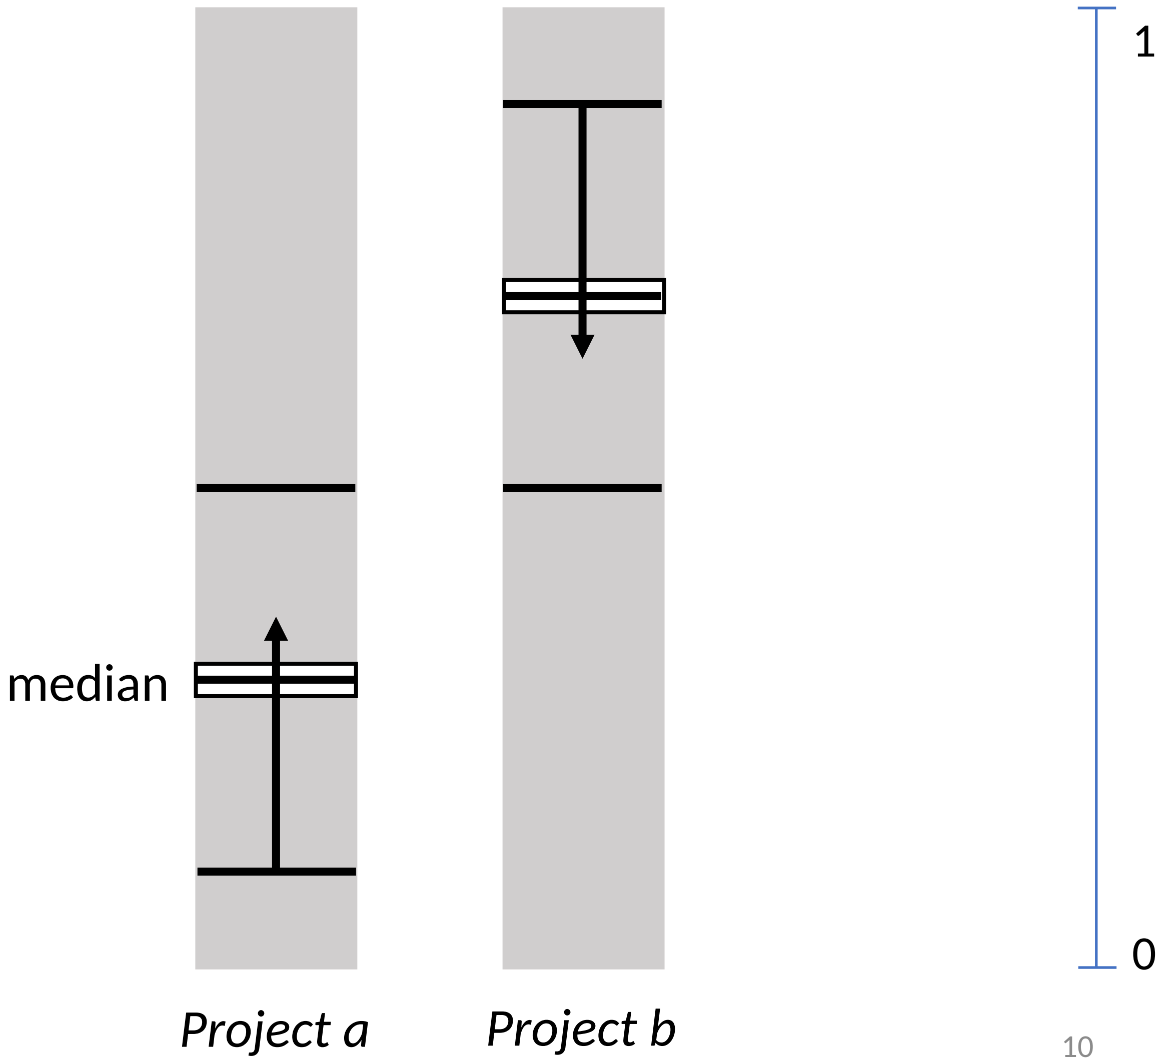

Two Projects – Generalized Medians

- One-dimensional: $p_b = 1 - p_a$
  - Aggregating n numbers over which voters have single-peaked preferences
- Moulin [1980] and follow-up work:
  Aggregator is anonymous, continuous, strategyproof
  if and only if it is a generalized median.
- Generalized median is defined by $n + 1$ phantom values $\alpha_0, \ldots, \alpha_n$
  and we return $q_a = \text{med}\{p_{1,a}, \ldots, p_{n,a}, \alpha_0, \ldots, \alpha_n\}$,

Two Projects – Generalized Medians

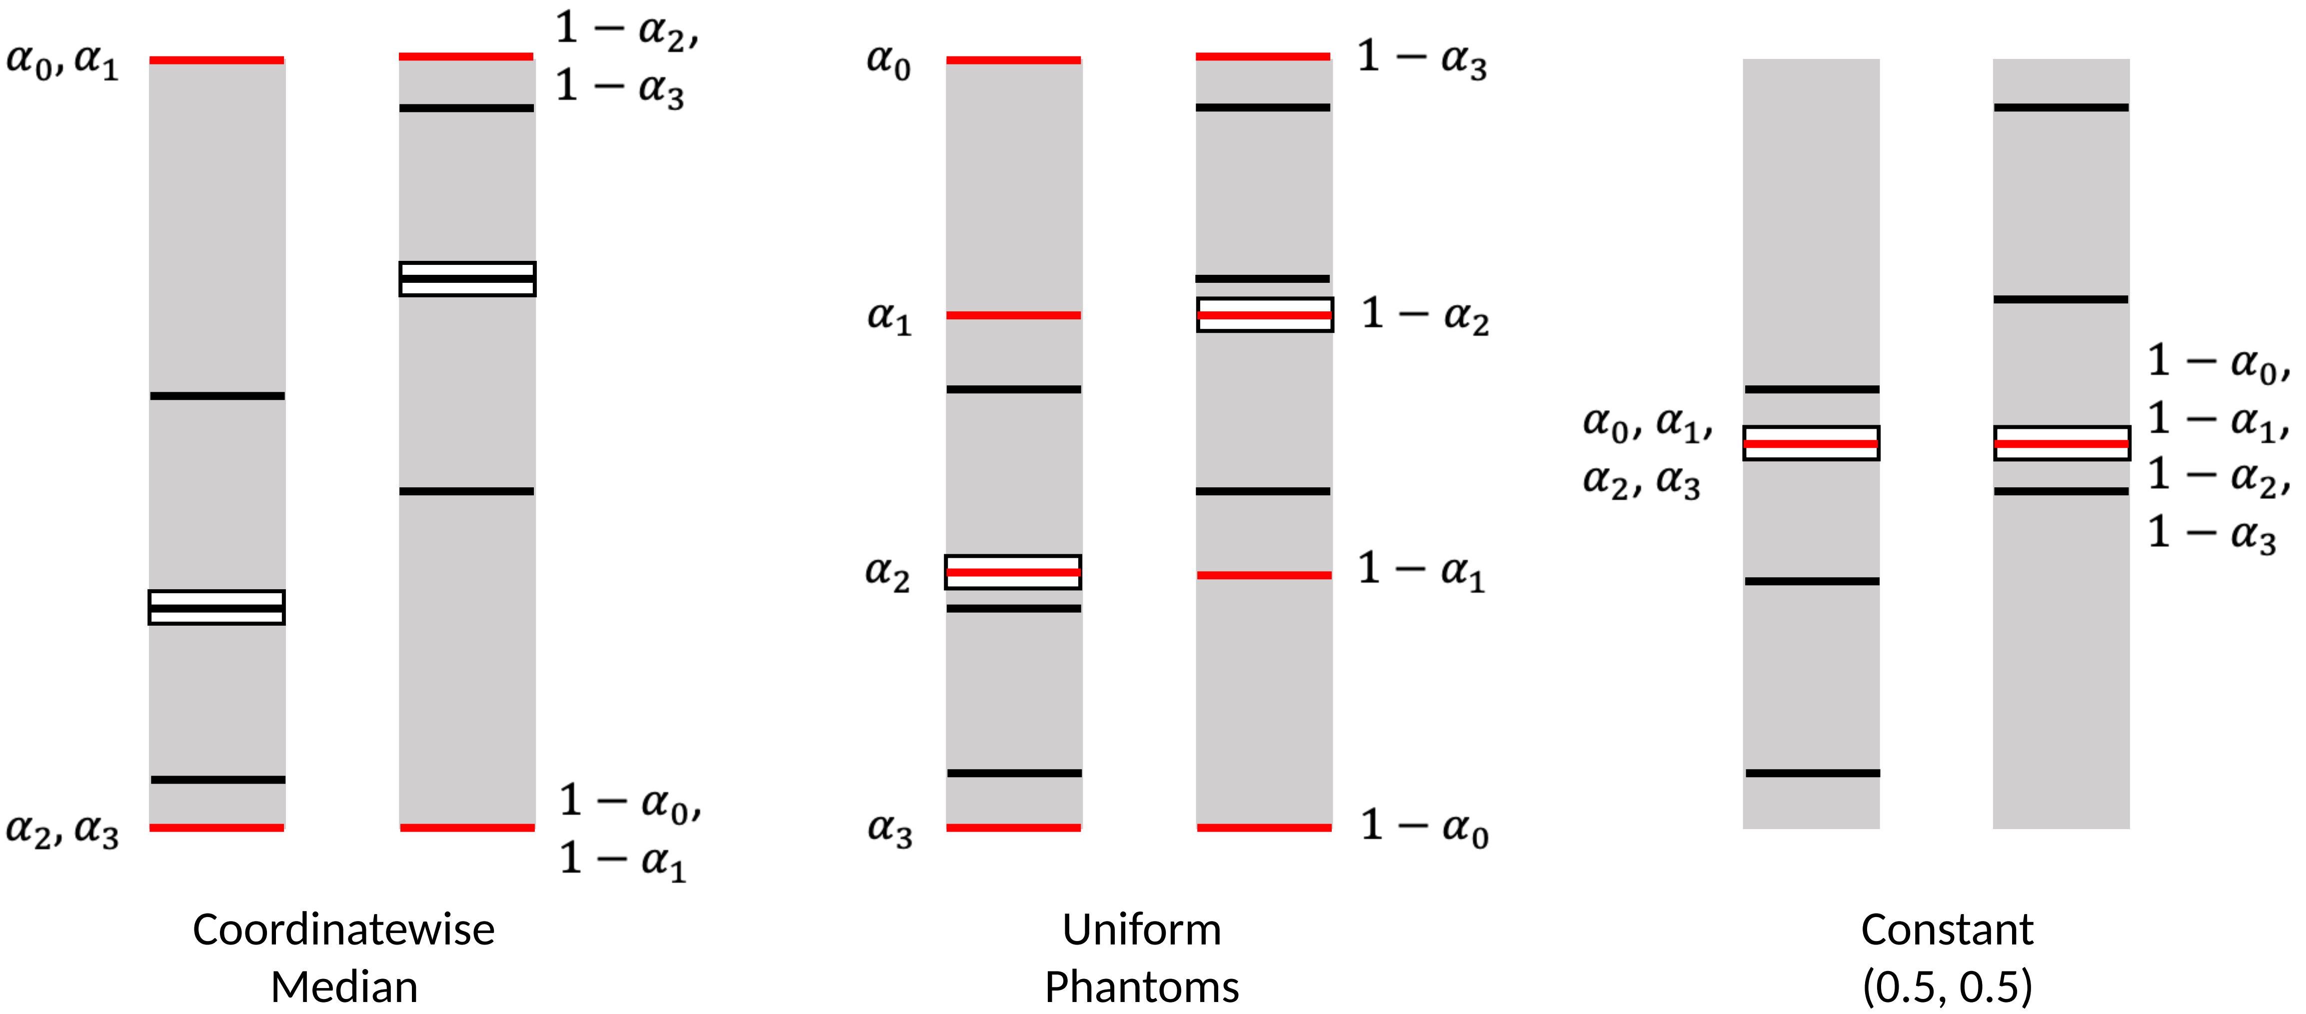

Two Projects – Generalized Medians

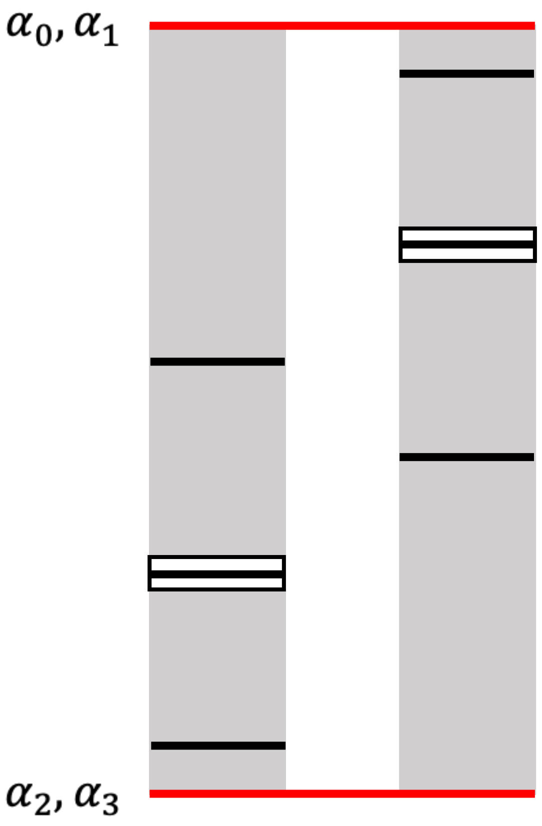

Coordinatewise
Median

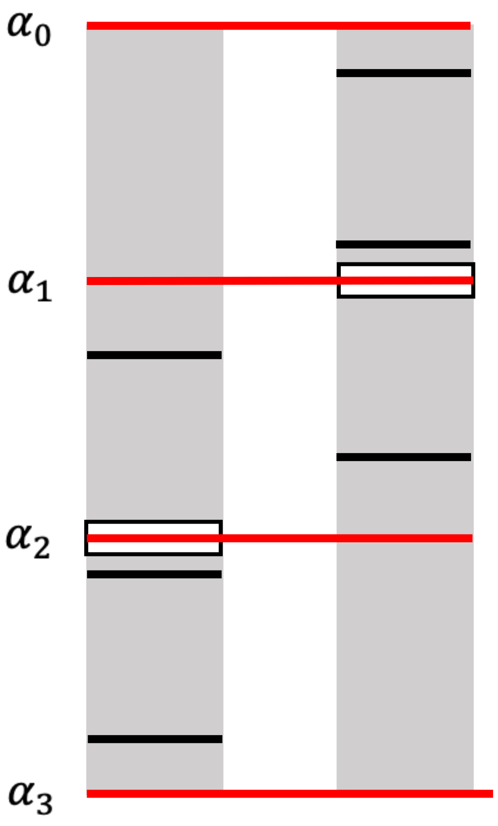

Uniform
Phantoms

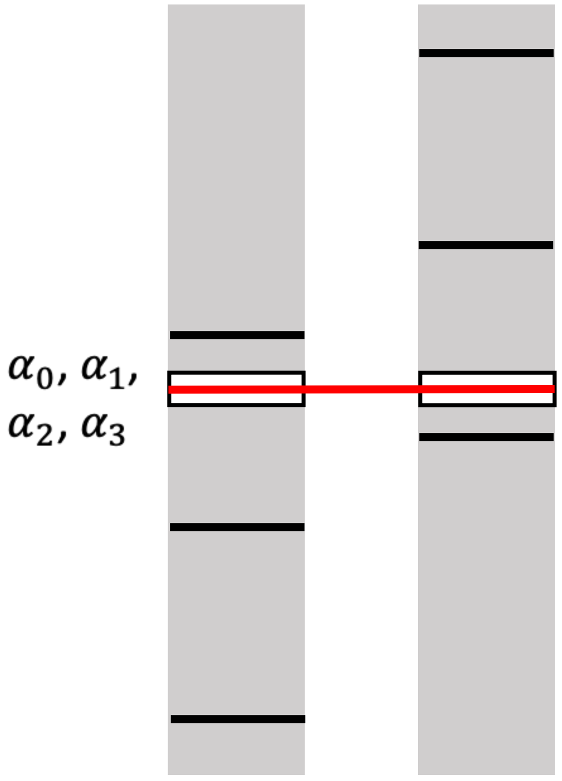

Constant
(0.5, 0.5)

Three or More Projects

What to do in higher dimensions?

- First idea: Take (generalized) medians on each coordinate
  - Fails to satisfy normalization constraint
- Multiplicative renormalization violates strategyproofness

$$\begin{matrix}(0, {}^{1}\!/_{4}, {}^{3}\!/_{4})\\ ({}^{3}\!/_{4}, 0, {}^{1}\!/_{4})\\ ({}^{1}\!/_{4}, {}^{3}\!/_{4}, 0)\end{matrix} \xrightarrow{\text{Median}} ({}^{1}\!/_{4}, {}^{1}\!/_{4}, {}^{1}\!/_{4}) \xrightarrow{\text{Normalize}} ({}^{1}\!/_{3}, {}^{1}\!/_{3}, {}^{1}\!/_{3})$$

$$\begin{matrix}(0, {}^{1}\!/_{6}, {}^{5}\!/_{6})\\ ({}^{3}\!/_{4}, 0, {}^{1}\!/_{4})\\ ({}^{1}\!/_{4}, {}^{3}\!/_{4}, 0)\end{matrix} \xrightarrow{\text{Median}} ({}^{1}\!/_{4}, {}^{1}\!/_{6}, {}^{1}\!/_{4}) \xrightarrow{\text{Normalize}} ({}^{3}\!/_{8}, {}^{1}\!/_{4}, {}^{3}\!/_{8})$$

Moving Phantom Mechanis

[F., Pennock, Peters, Wortman Vaughan 20

- Idea:
  - If sum of medians is > 1, move phantoms down
  - If sum of medians is < 1, move phantoms up
  - Converge to = 1
- A moving phantom mechanism is given by a “movie” $f\colon [0,1] \to [0,1]^{n+1}$ with $f(0) = \mathbf{0}$ and $f(1) = \mathbf{1}$, continuous, increasing.
- We play the movie until $f(t^*)$ gives phantoms that normalize output.
- Intermediate Value Theorem: t^* exists

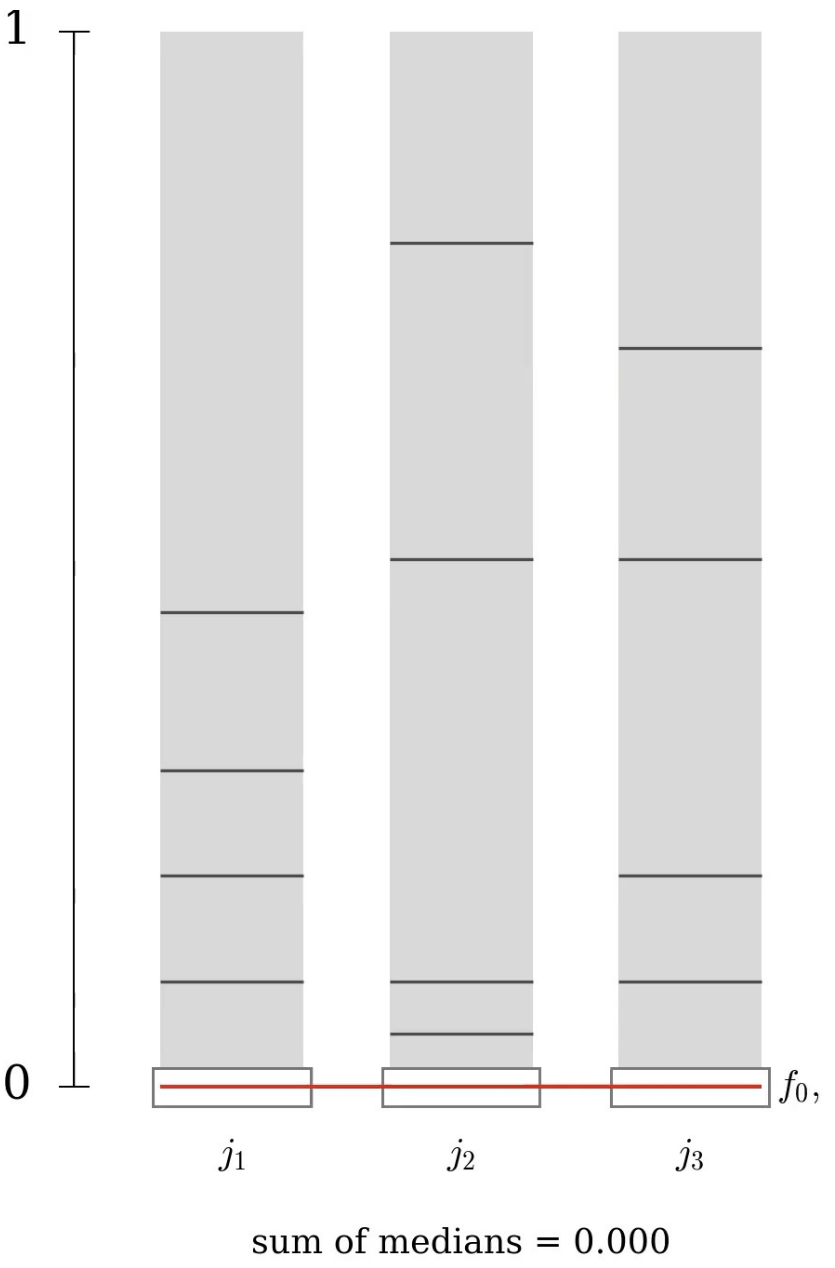

Moving Phantoms are Strategyproof

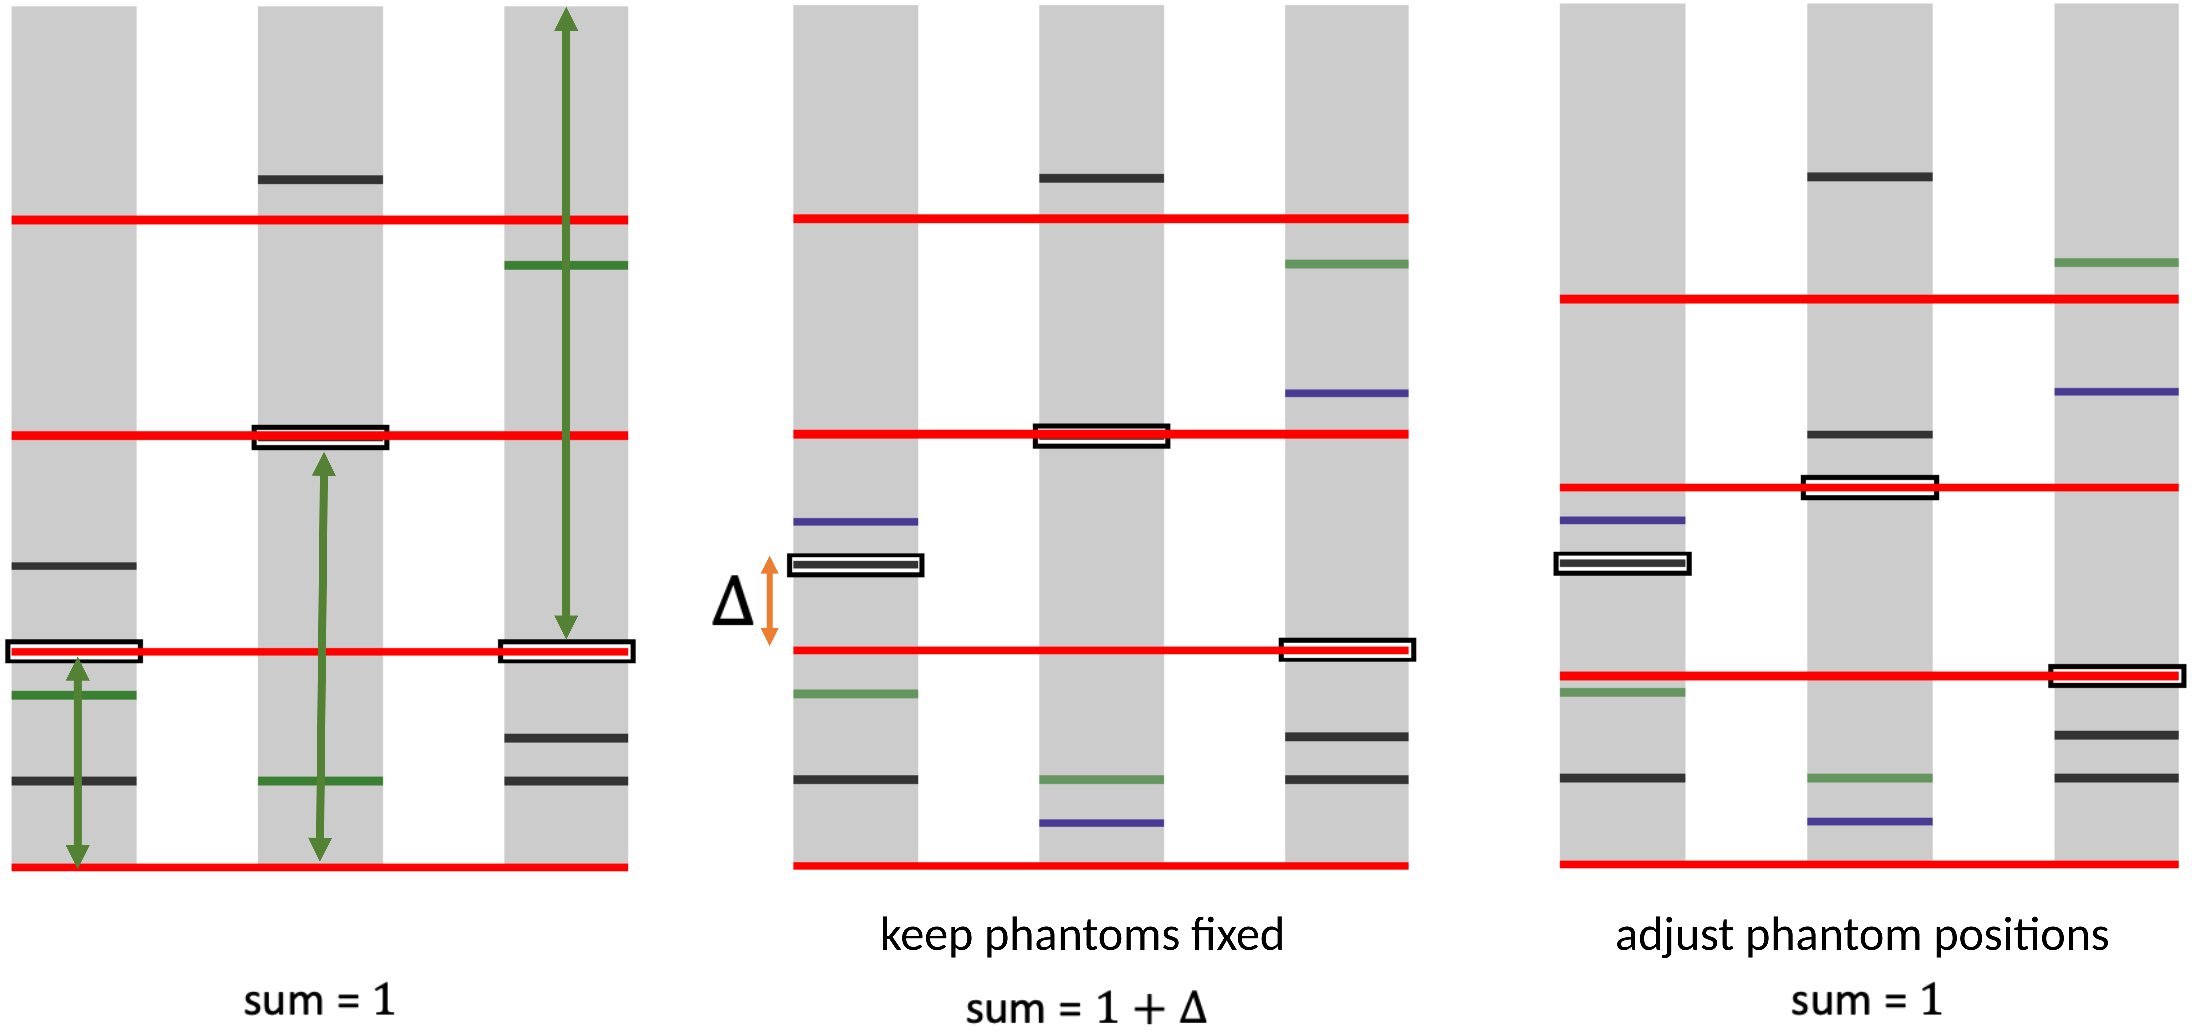

Outline

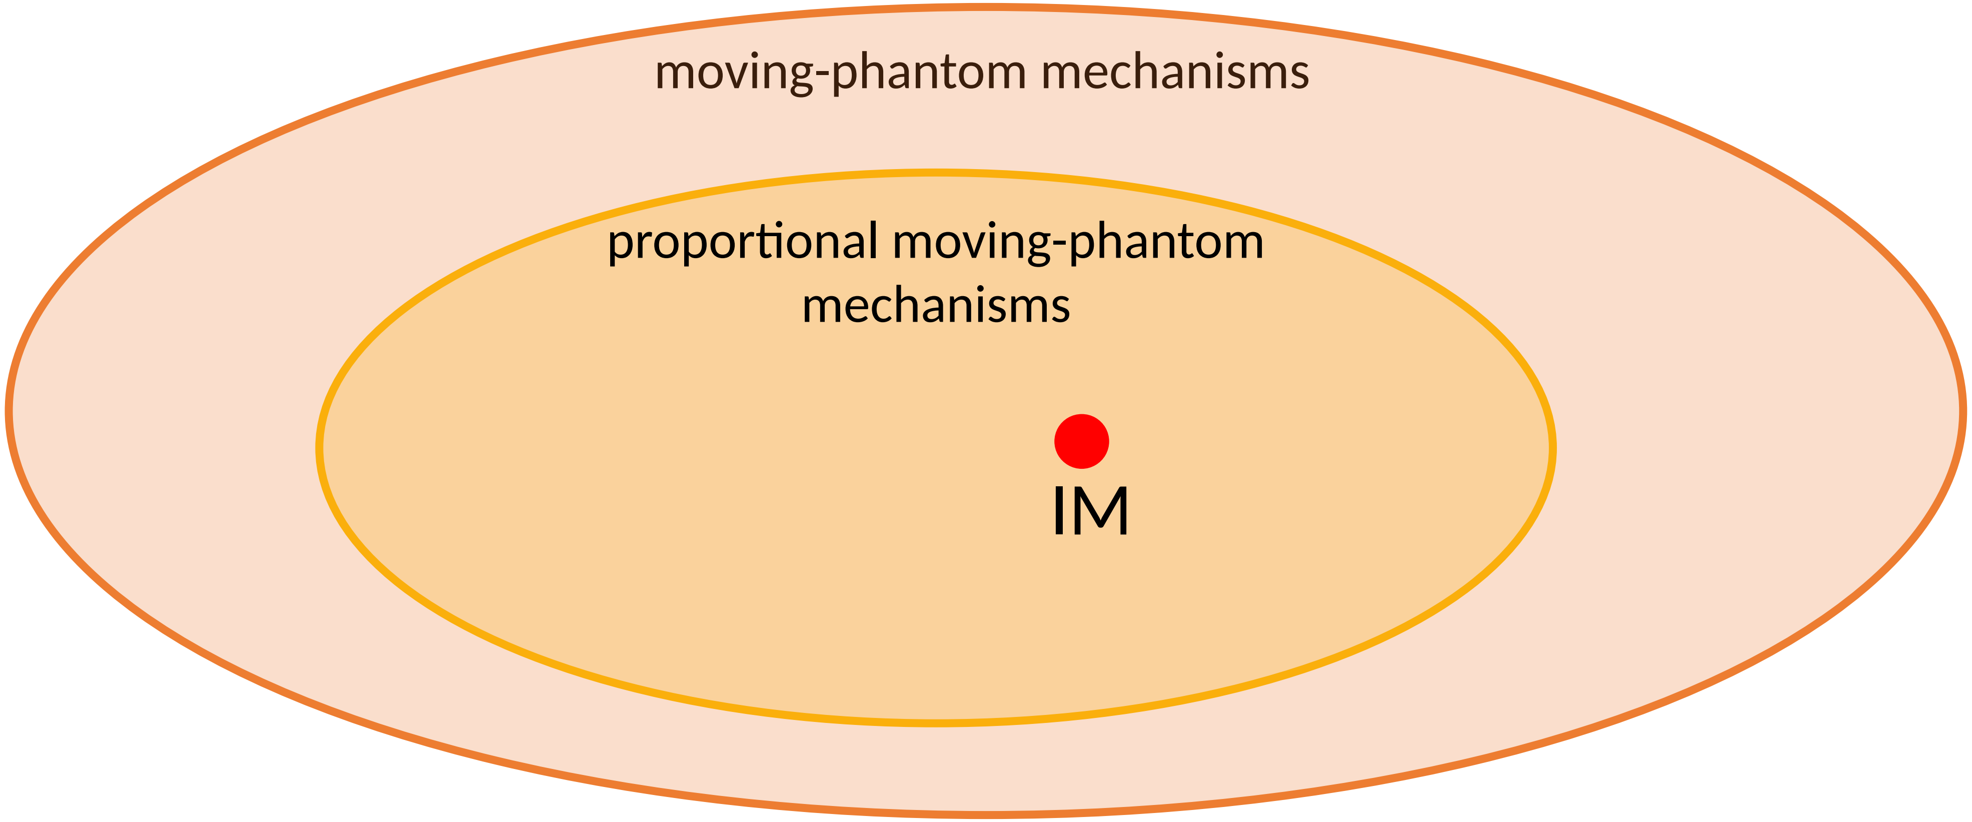
moving-phantom mechanisms
proportional moving-phantom
mechanisms
IM

Outline

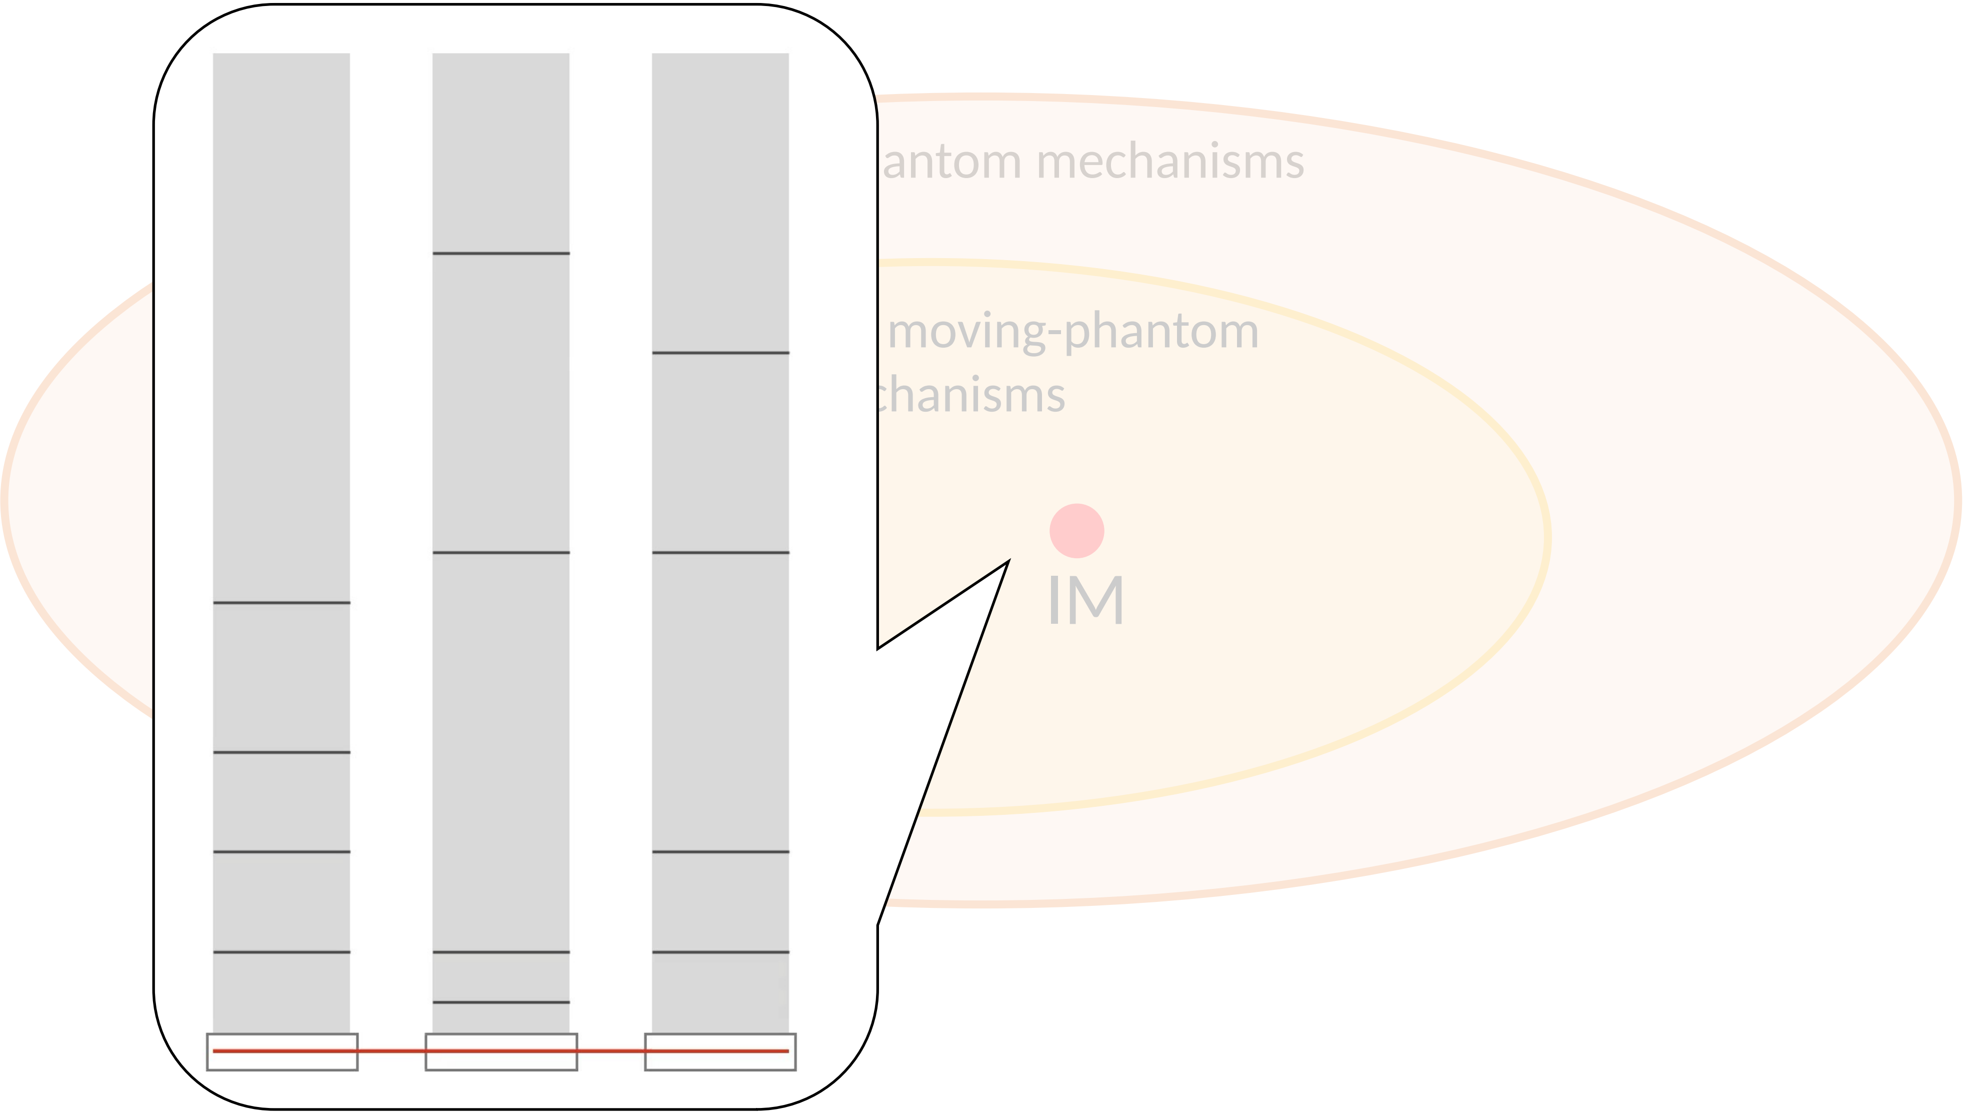

Independent Markets (IM) mechanism

Place phantoms uniformly between 0 and rising y.

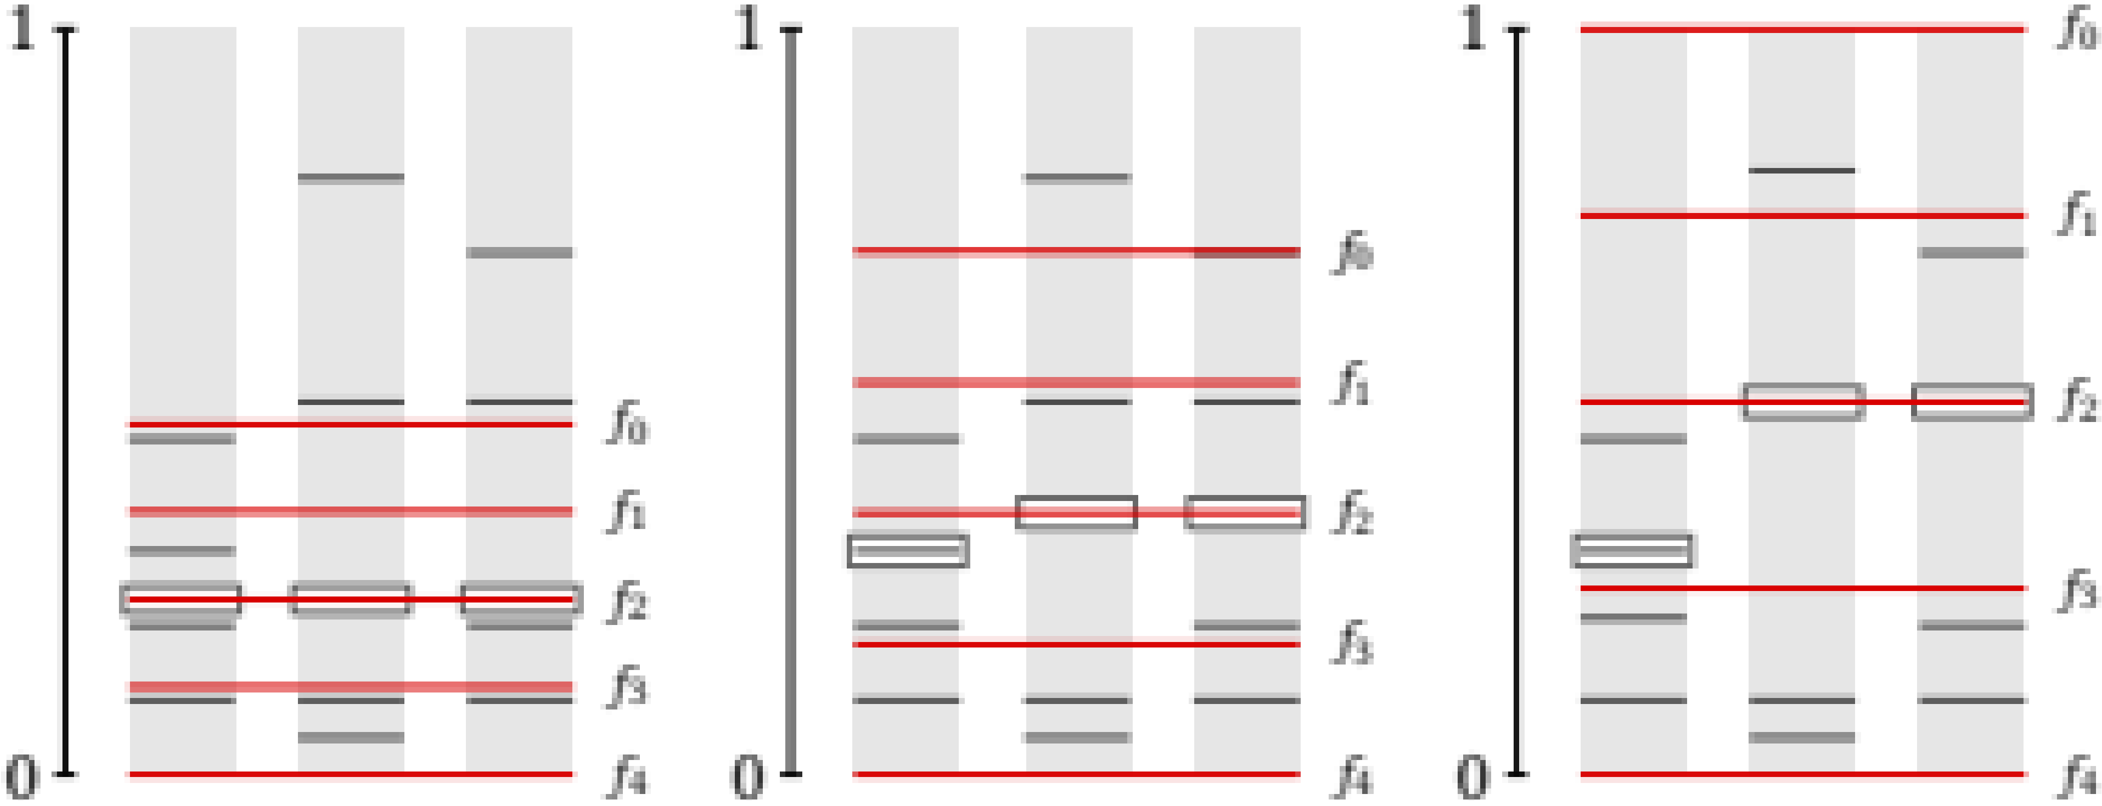

Market Clearing Price with One Commodity

- n buyers with budget \$1 each
- Supply: x divisible units of a good
- each buyer i has a valuation v_i per unit of the good
- at price π, buyer i demands
  - ∞ units if $\pi = 0$
  - $1/\pi$ units if $v_i > \pi$
  - 0 units if $v_i \leq \pi$
- Clearing price $\pi^* = \sup \{\pi : \sum_{i \in [n]} D_i(\pi) > x\}$

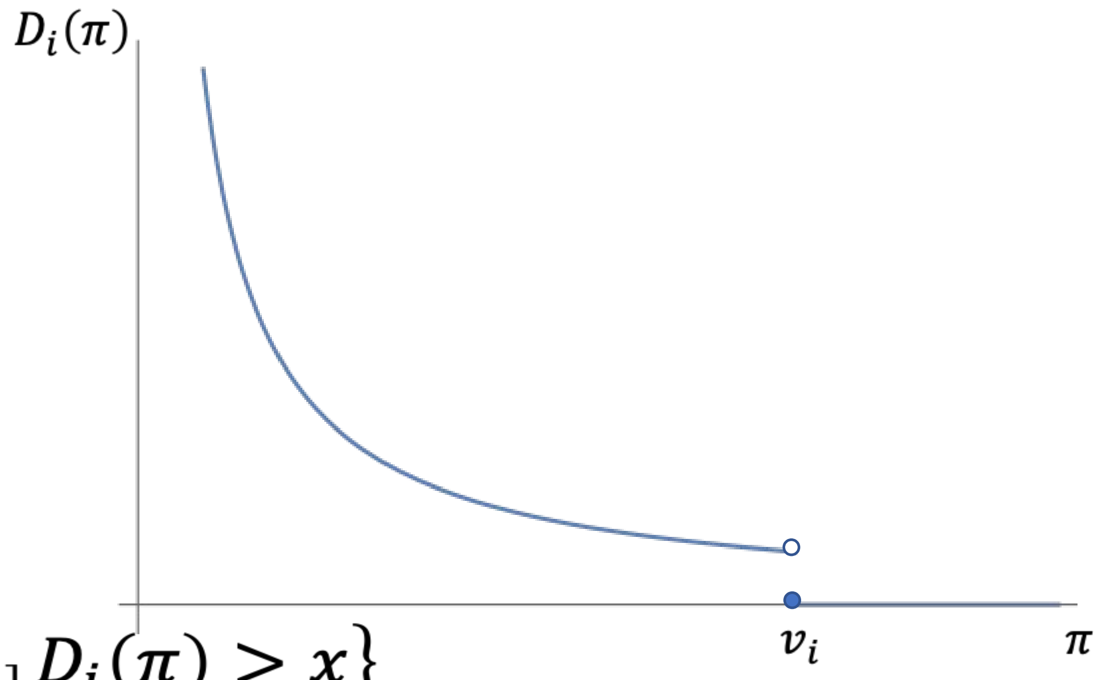

Theorem: We have $\pi^* = \text{med}\left(v_1, \ldots, v_n, 0, \frac{1}{x}, \frac{2}{x}, \ldots, \frac{n-1}{x}, \frac{n}{x}\right)$.

Market Clearing Price with One Commodity

$x = 2$

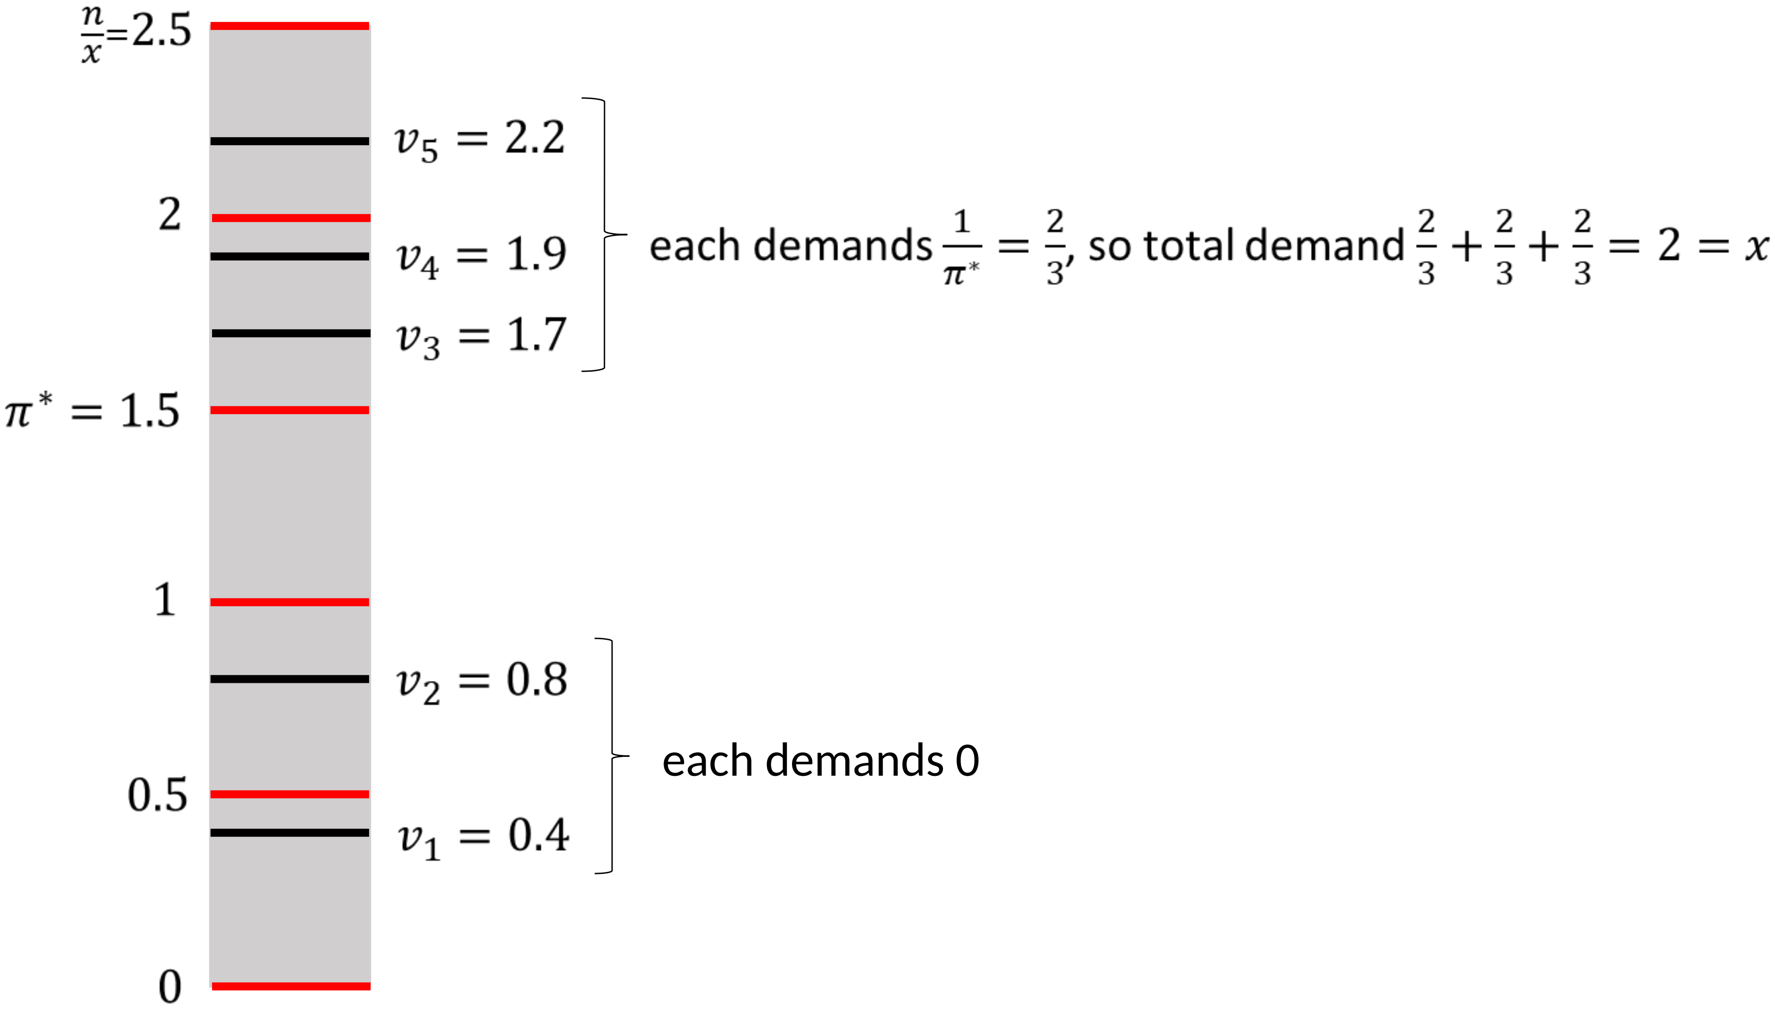

Independent Markets (IM) mechanism

Place phantoms uniformly between 0 and rising y.

Equivalently:

- Set up m independent markets, each selling $x = n/y$ units.
- Each voter has $1 in each market. Valuation = report.
- Mechanism varies x until the clearing prices sum to 1.

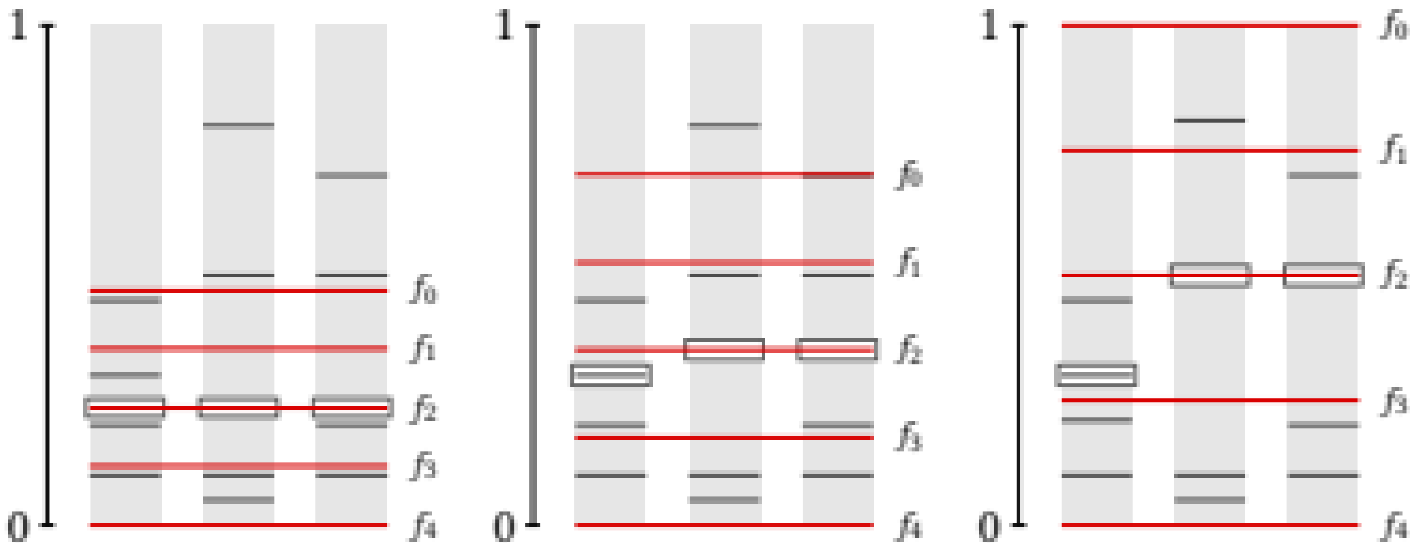

Voting Game

- Every voter i selects a spending level $s_{i,x} \in [0,1]$ for every project x
- Projects are funded in proportion to the amount spent on them
- **Theorem**: The independent markets outcome is the same as the outcome of the game at equilibrium.

	Project a	Project b	Project c
Voter 1	1	0.8	0
Voter 2	0.6	0	0.3
Voter 3	0.4	0.2	0.7
	2.0	1.0	1.0
	0.5	**0.25**	**0.25**

Single-Minded Proportionality

- Theorem: A moving-phantom mechanism is single-minded proportional if, and only if, the phantom system "passes through" uniform positions.

- Corollary: Independent Markets satisfies proportionality.

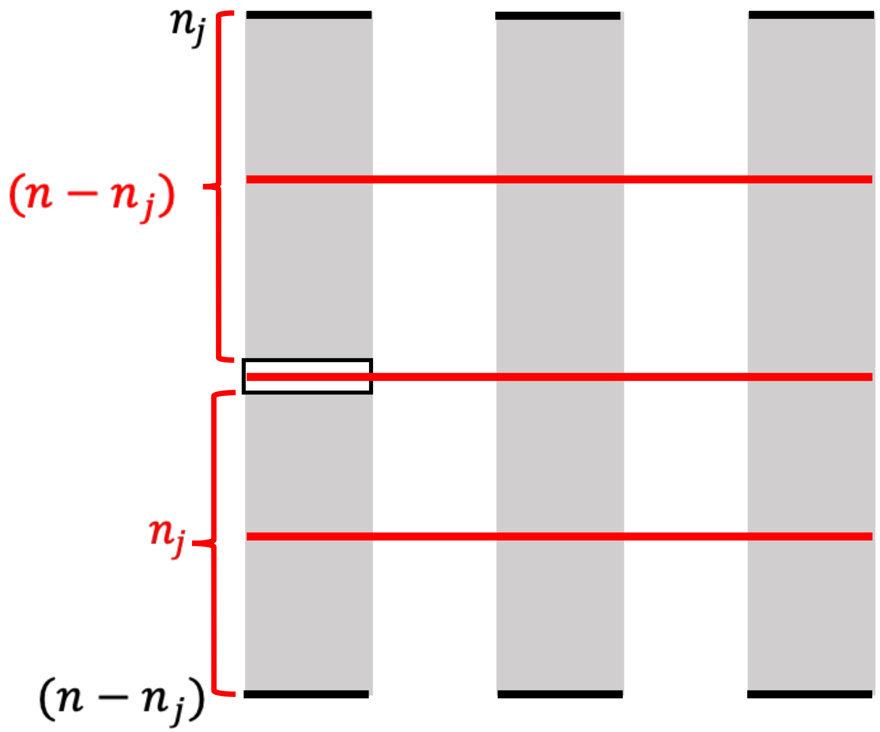

A Bad Example

- $p_1 = (0.8, 0.2, 0)$
  $p_2 = (0.8, 0, 0.2)$

- Output: $(0.6, 0.2, 0.2)$

- But both voters would prefer $(0.8, 0.1, 0.1)$
  - Violation of Pareto optimality

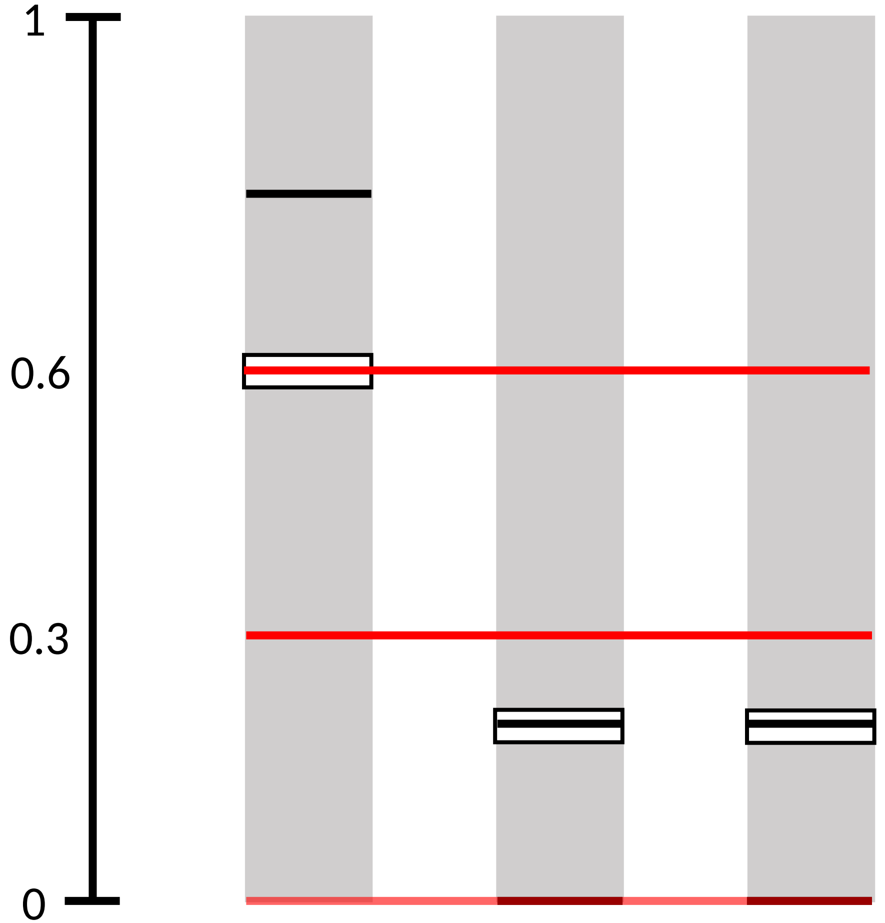

Outline

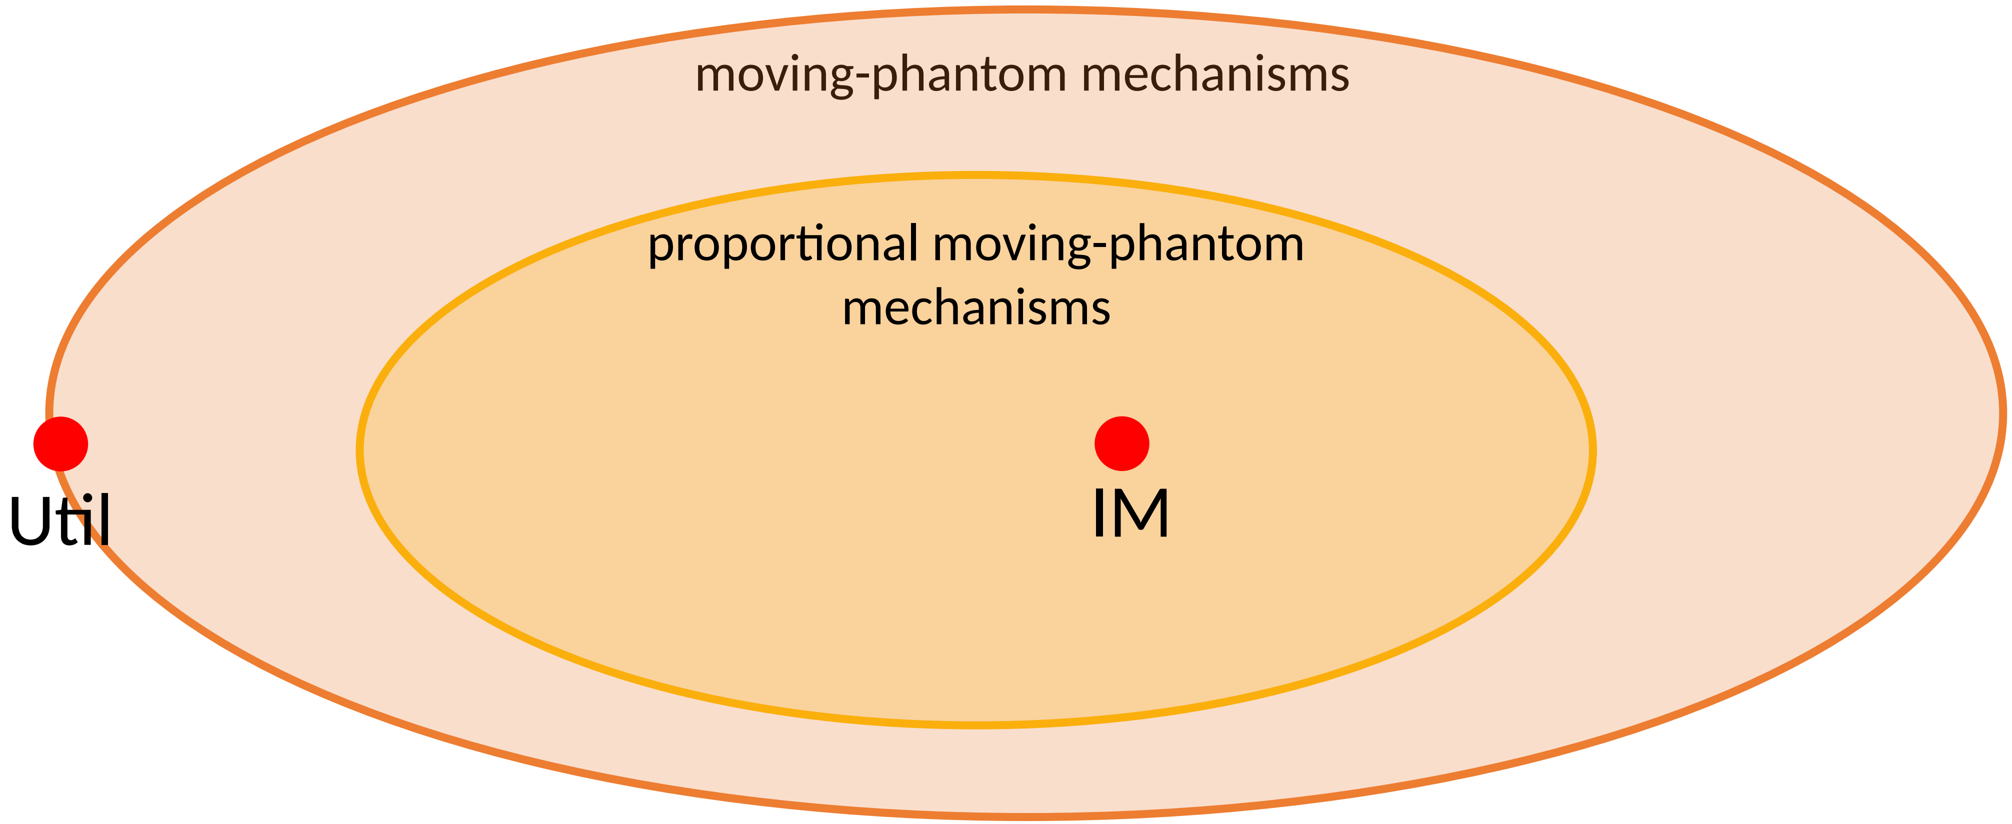
moving-phantom mechanisms
proportional moving-phantom mechanisms
Util
IM

Outline

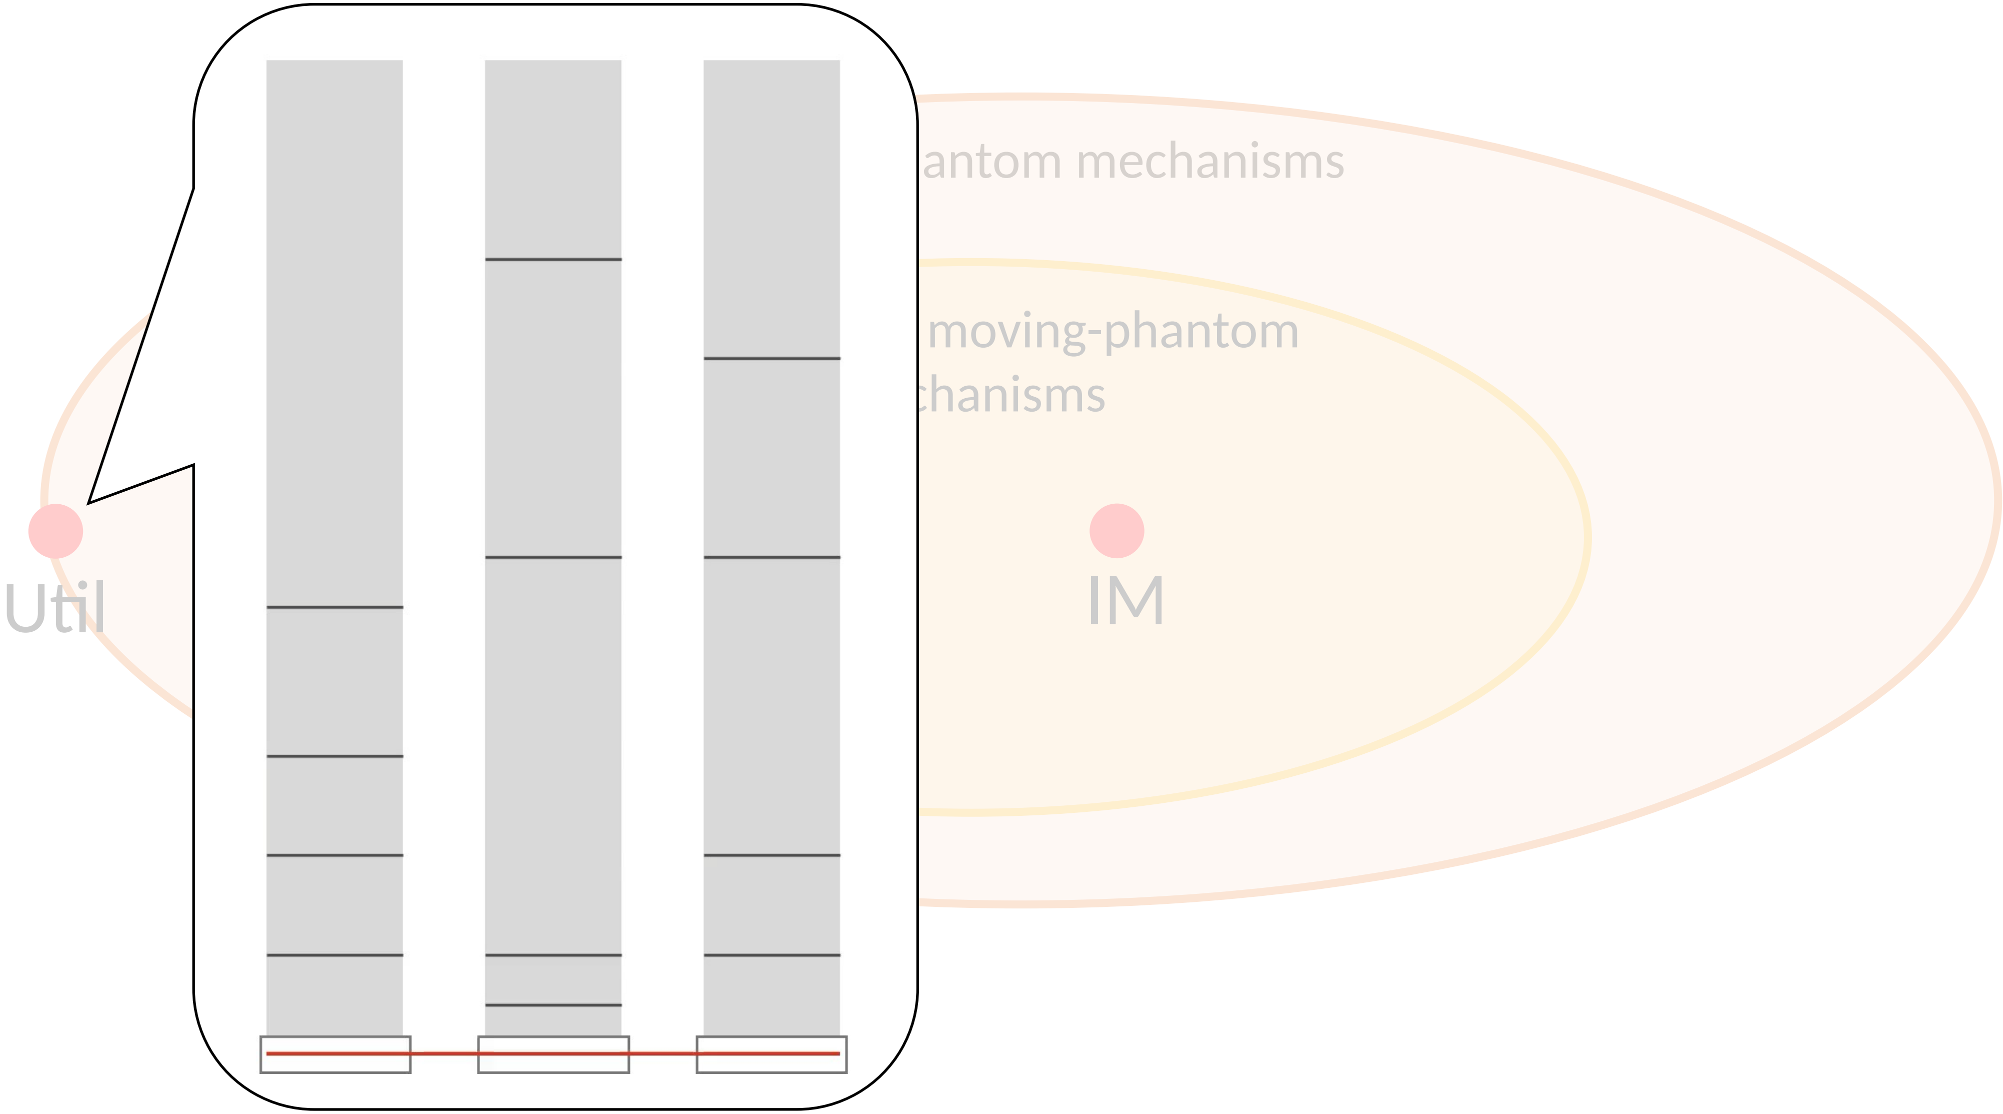

Pareto Optimality

- A distribution p is Pareto dominated by q if $d_i(p_i, p) \geq d_i(p_i, q)$ for all i & $d_i(p_i, p) > d_i(p_i, q)$ for some i.
- A mechanism is Pareto optimal if its output is never dominated.
- Moving phantom mechanisms are not usually Pareto optimal.

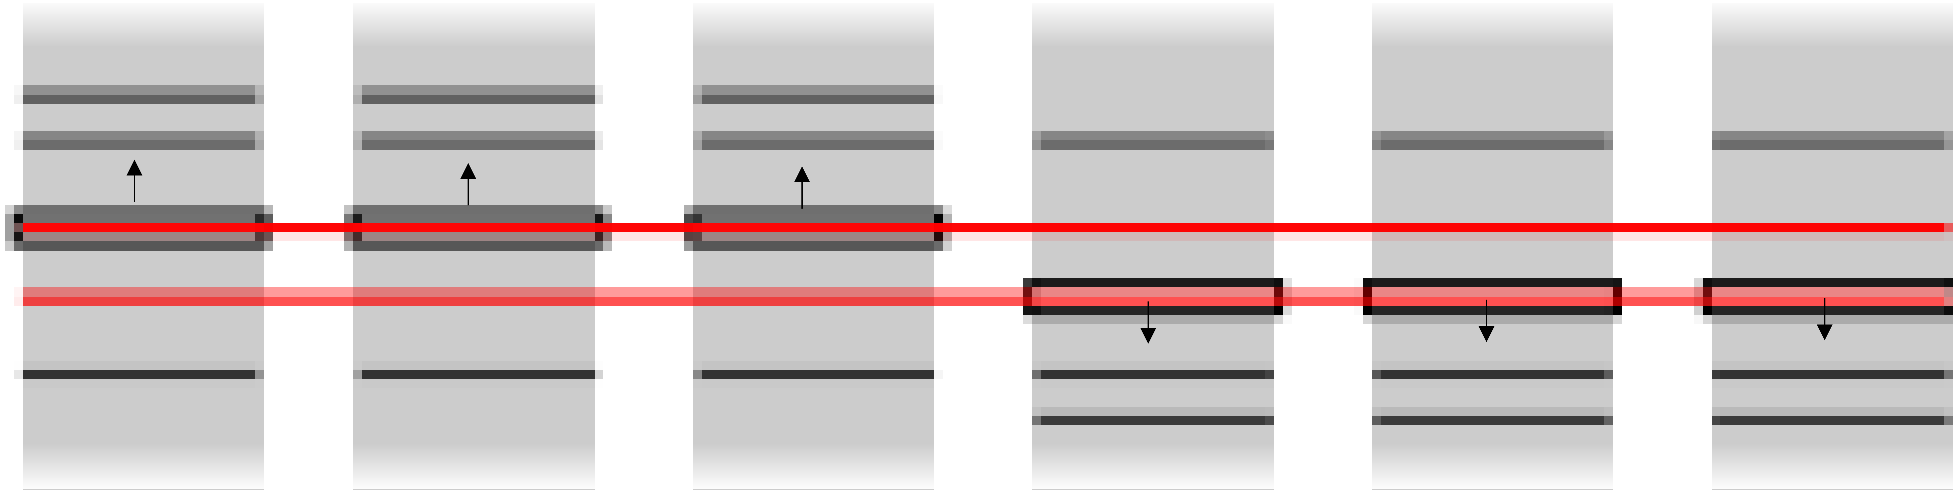

Pareto Optimality: Characterization

- **Theorem:**
  - If at some t, there are two interior phantoms with $0 < f_i, f_j < 1$, the resulting mechanism fails Pareto optimality at some profile.
  - There is a unique phantom movement that avoids having two interior phantoms.
  - The resulting mechanism is Pareto optimal.
  - The mechanism selects the distribution that:
    - maximizes utilitarian ℓ_1 social welfare
    - breaks ties in favor of higher (Shannon) entropy

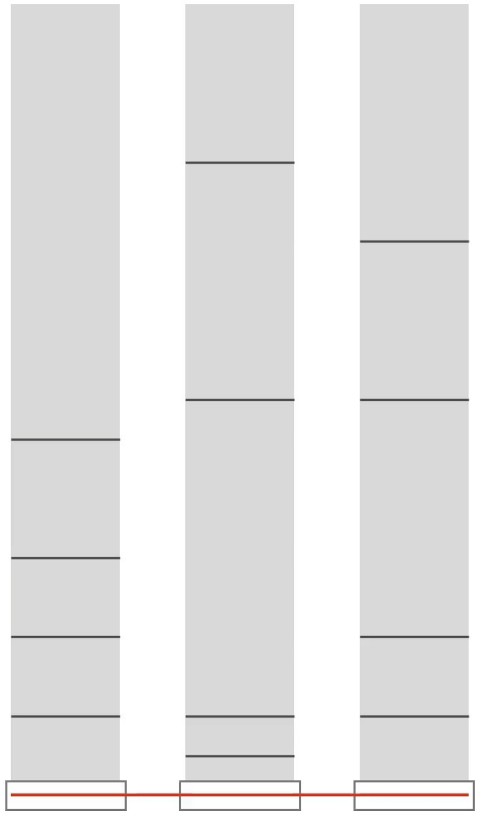

Example

- $p_1 = (0.8, 0.2, 0)$
  $p_2 = (0.8, 0, 0.2)$

- Output: $(0.8, 0.1, 0.1)$

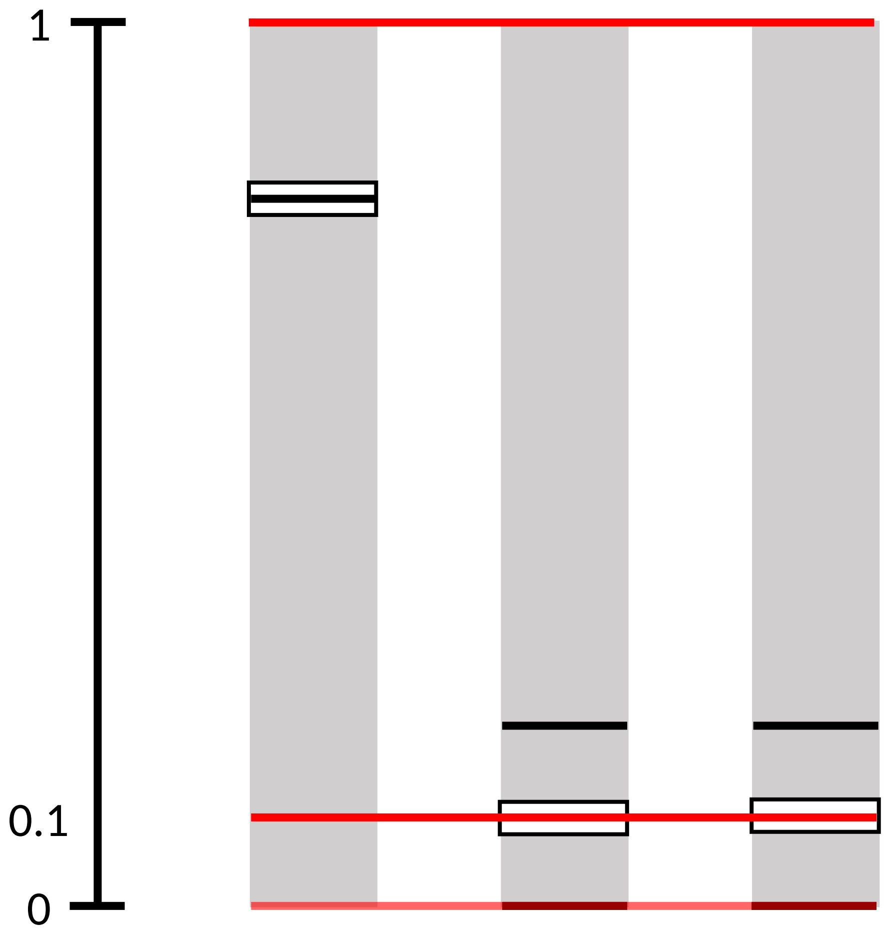

Outline

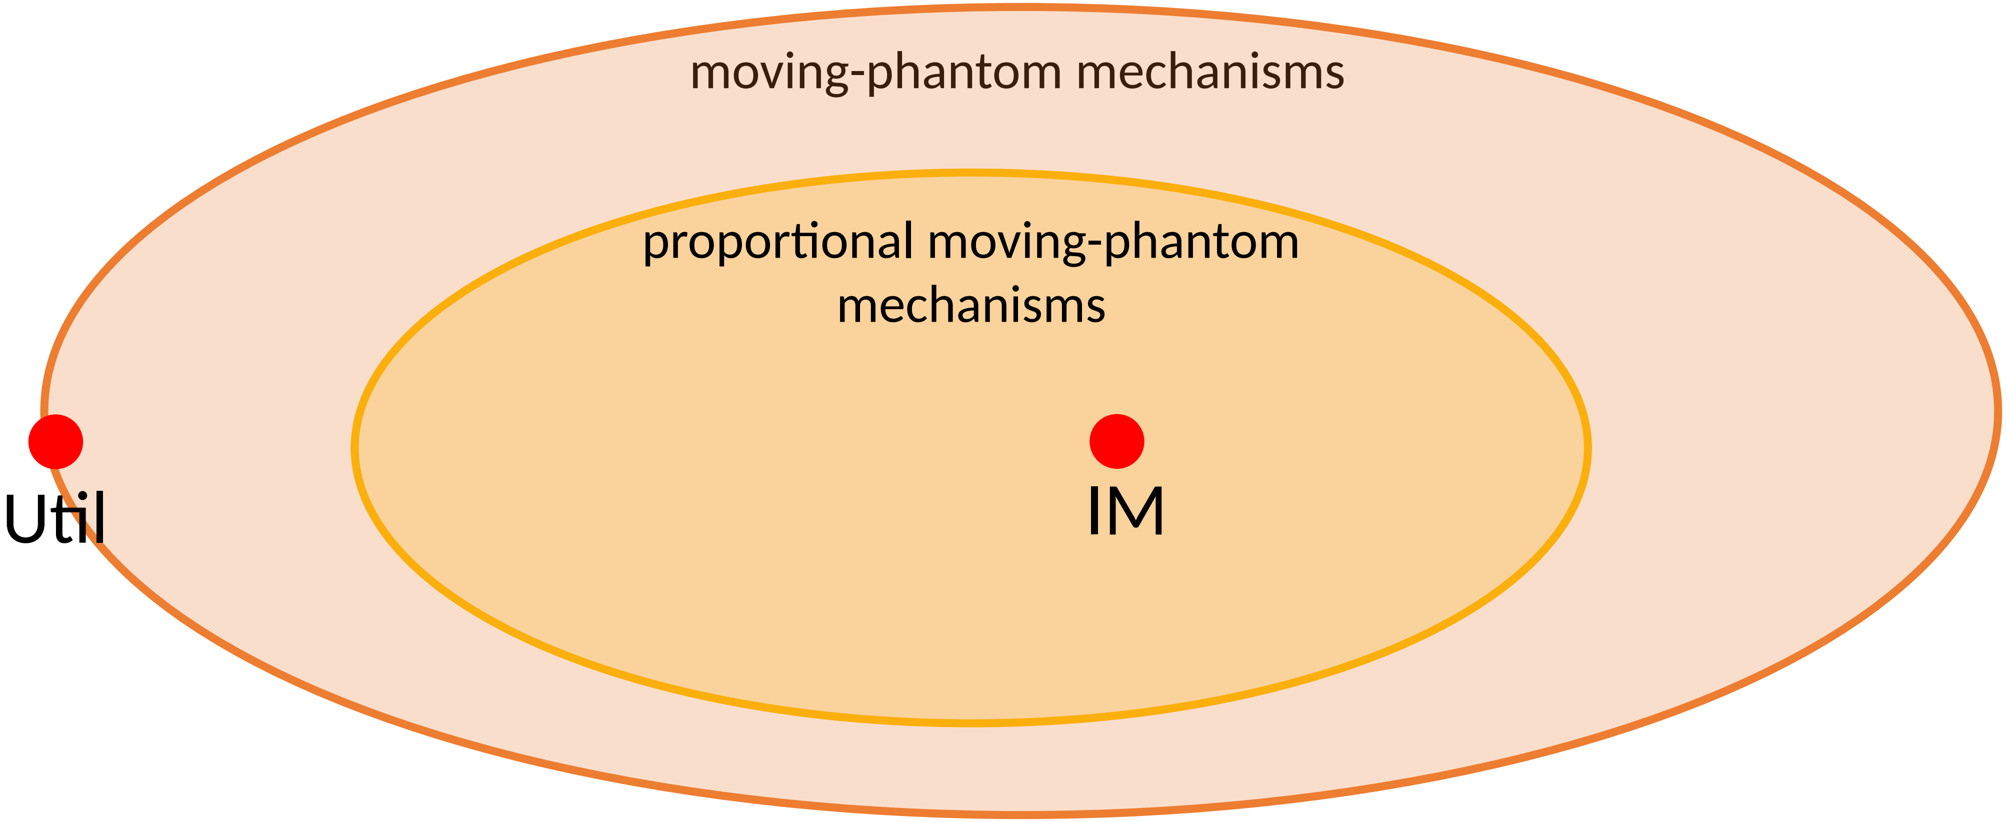
moving-phantom mechanisms
proportional moving-phantom
mechanisms
Util
IM

Outline

(continuous, anonymous, neutral) strategyproof mechanisms

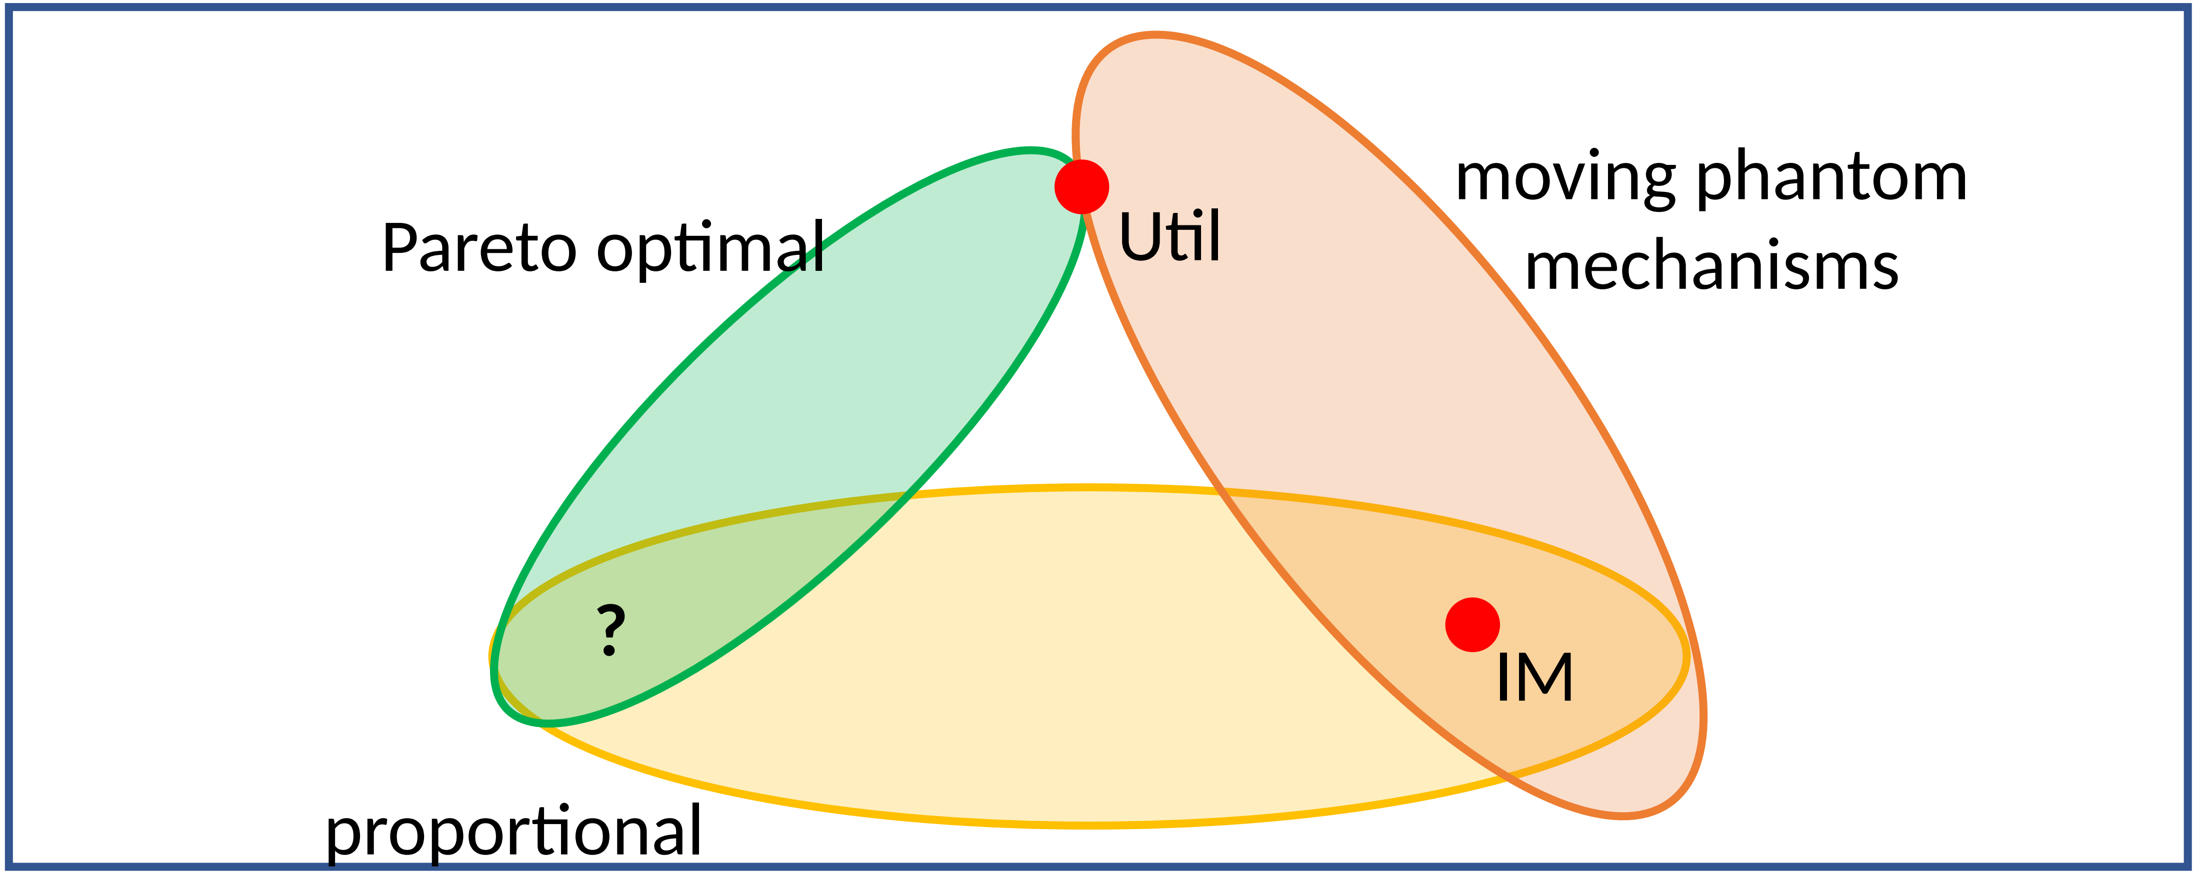

Outline

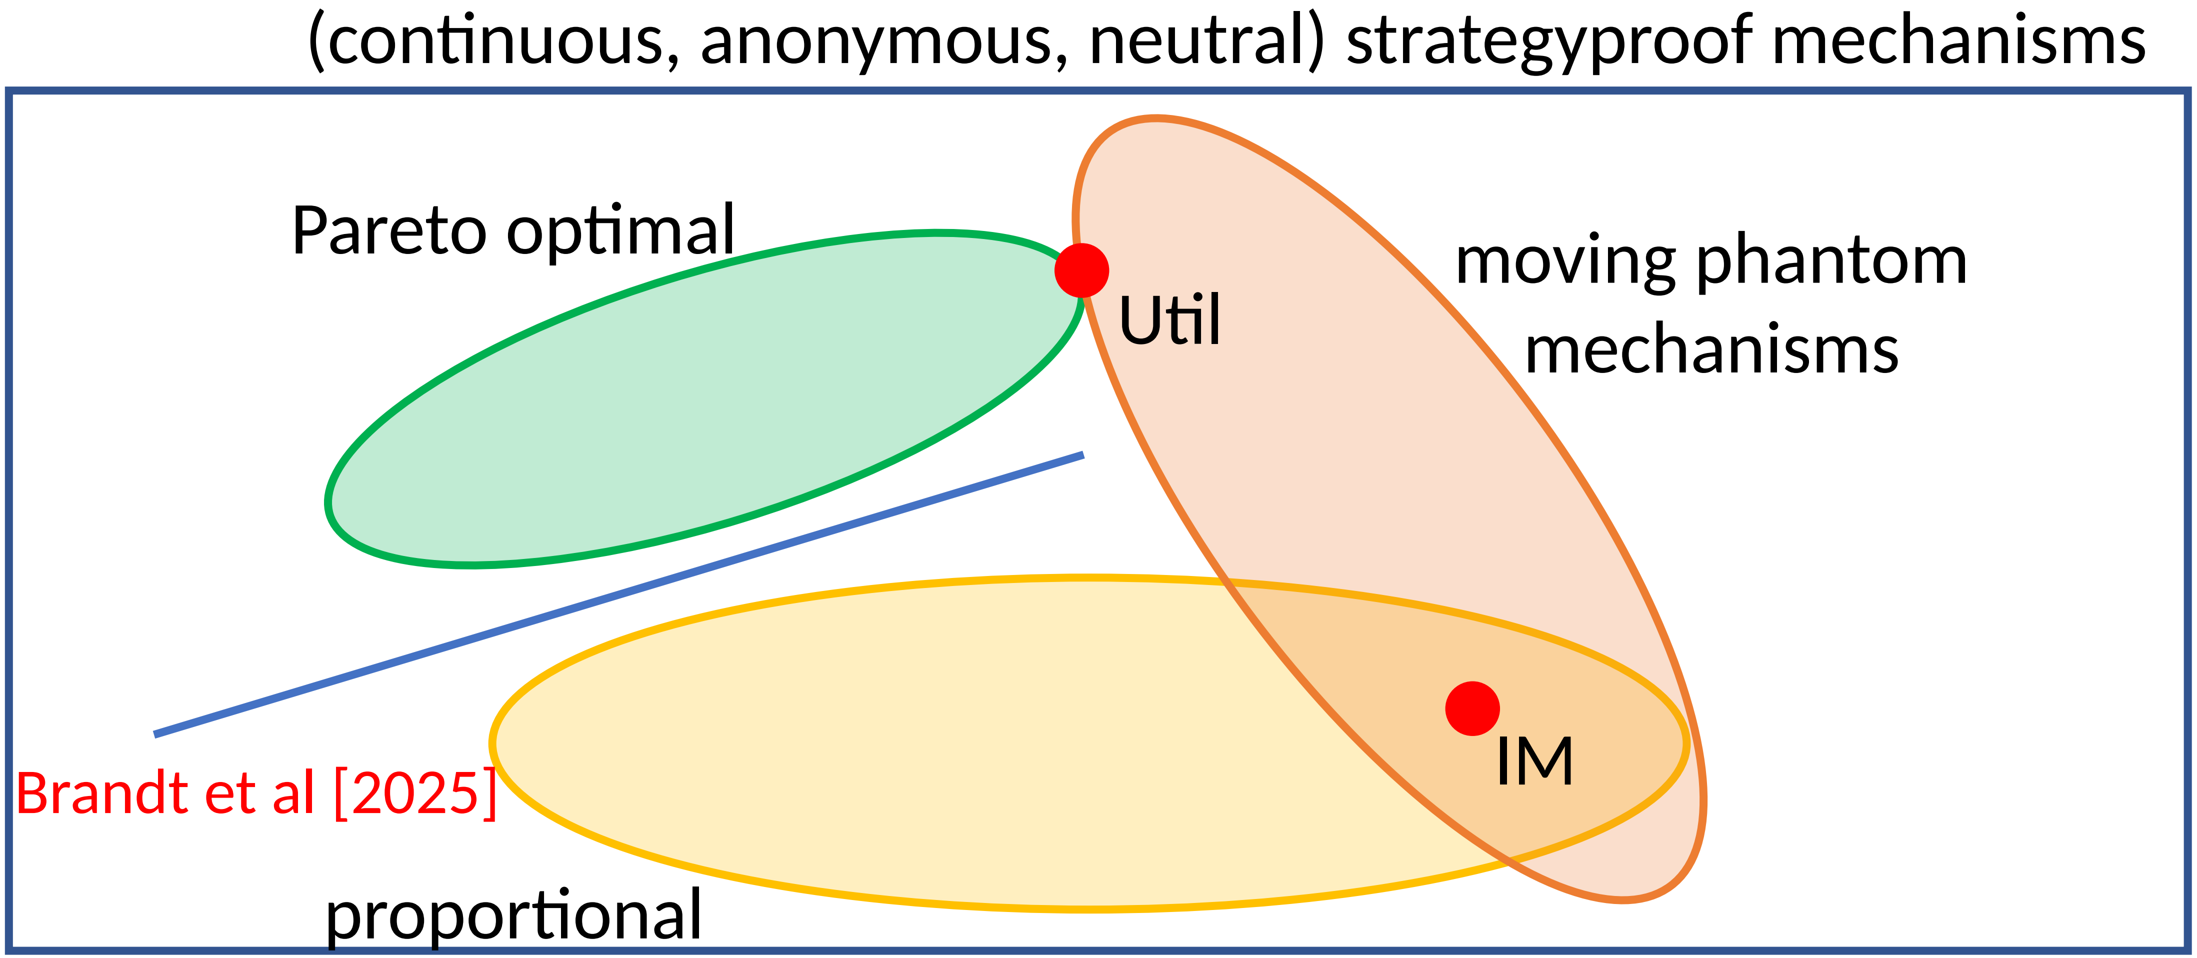
(continuous, anonymous, neutral) strategyproof mechanisms
Pareto optimal
Util
moving phantom mechanisms
Brandt et al [2025]
IM
proportional

Outline

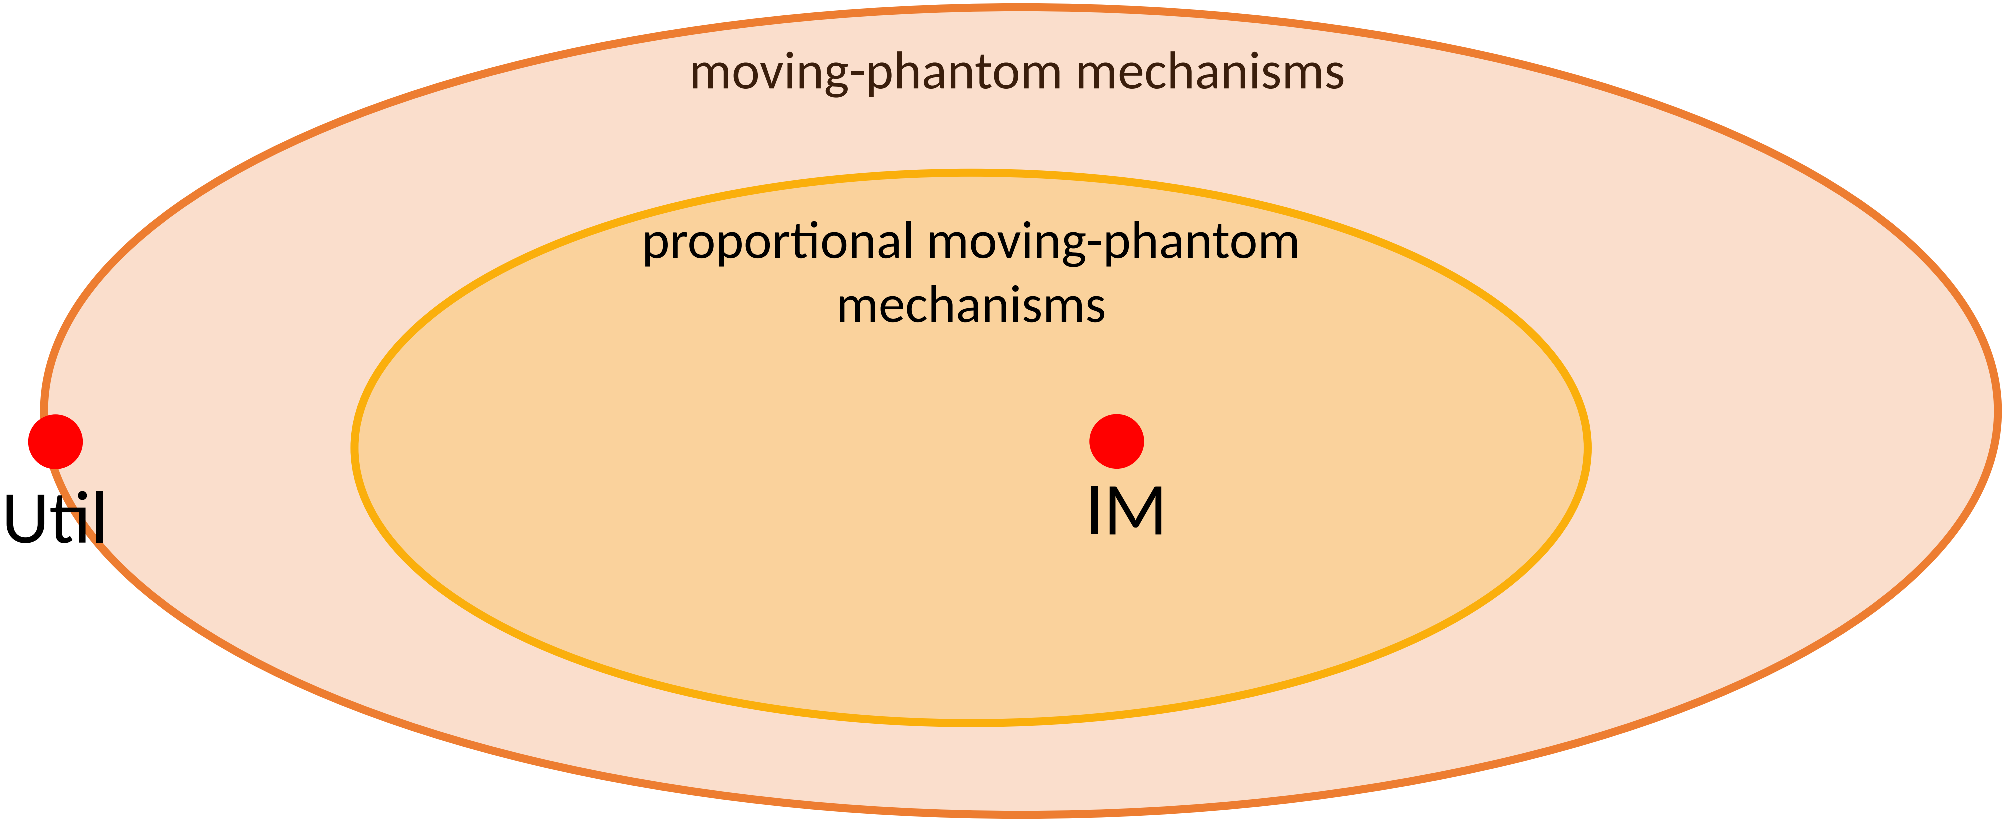
moving-phantom mechanisms
proportional moving-phantom mechanisms
Util
IM

Beyond Single-Minded Proportionality

- Single-minded proportionality is weak – only applies to single-minded profiles.
- Caragiannis et al [2022] propose evaluating mechanisms by the quality of their mean approximation.
  - Mechanism A is α-approximate if $d(A(P), \bar{P}) \leq \alpha$ for all preference profiles P.
- For $m = 2$, uniform phantoms achieve optimal mean approximation, $\alpha = \frac{1}{2}$ (cf. $\alpha = 1$ for standard median).

Mean Approximation Lower Bound

- Caragiannis et al [2022]: No moving-phantom mechanism is α-approximate for $\alpha < 1 - \frac{1}{m}$.

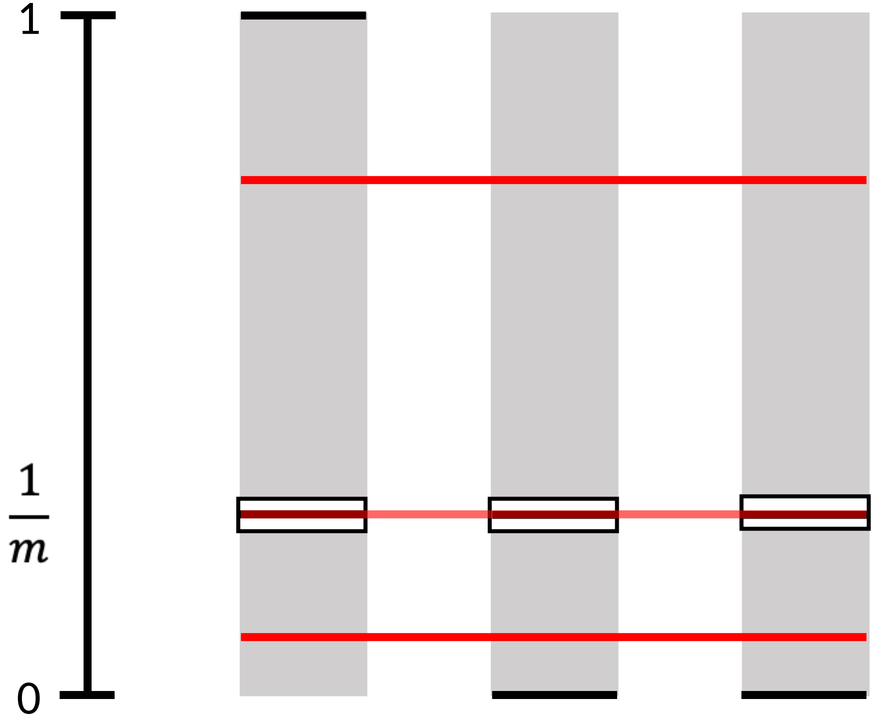

Piecewise Uniform Mechani

- Caragiannis et al [2022] define the Piecewise Uniform moving phantom mechanism, which is $(\frac{2}{3}+\epsilon)$-approximate for $m=3$.
  - $\epsilon < 10^{-5}$ gap comes from solving a non-linear program

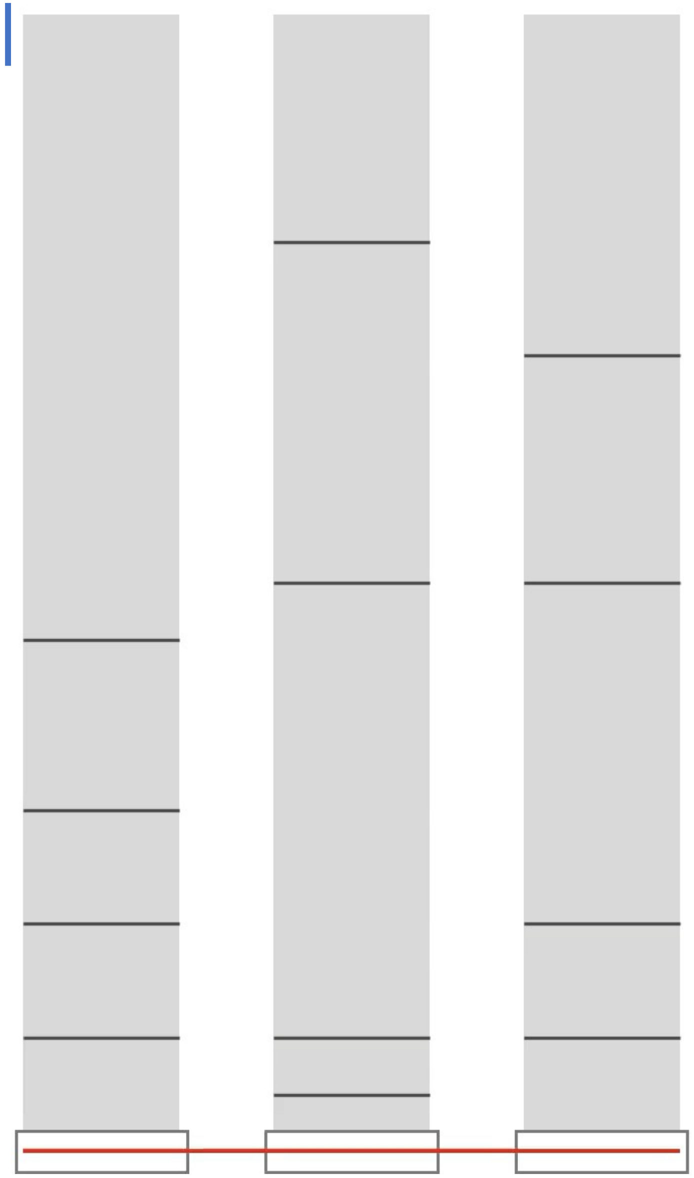

Project Fairness [Freeman and Schmidt-Kraepelin 2024]

- Can use other distances to mean, e.g., $\ell_\infty(p, q) = \max_{j \in [m]} |p_j - q_j|$
- Mechanism A is α-project fair if $\ell_\infty(A(P), \bar{P}) \leq \alpha$ for all preference profiles P.
  - "No project's funding is more than α from what they would receive under the mean allocation."
- No moving-phantom mechanism is α-project fair for $\alpha < \frac{1}{2}\left(1 - \frac{1}{m}\right)$.

Ladder Mechanism

- Ladder: Move phantoms one at a time, with phantom k starting to move once phantom $k-1$ reaches $1/n$.

- Theorem: Ladder is $\frac{1}{2}\left(1-\frac{1}{m}\right)$-project fair.

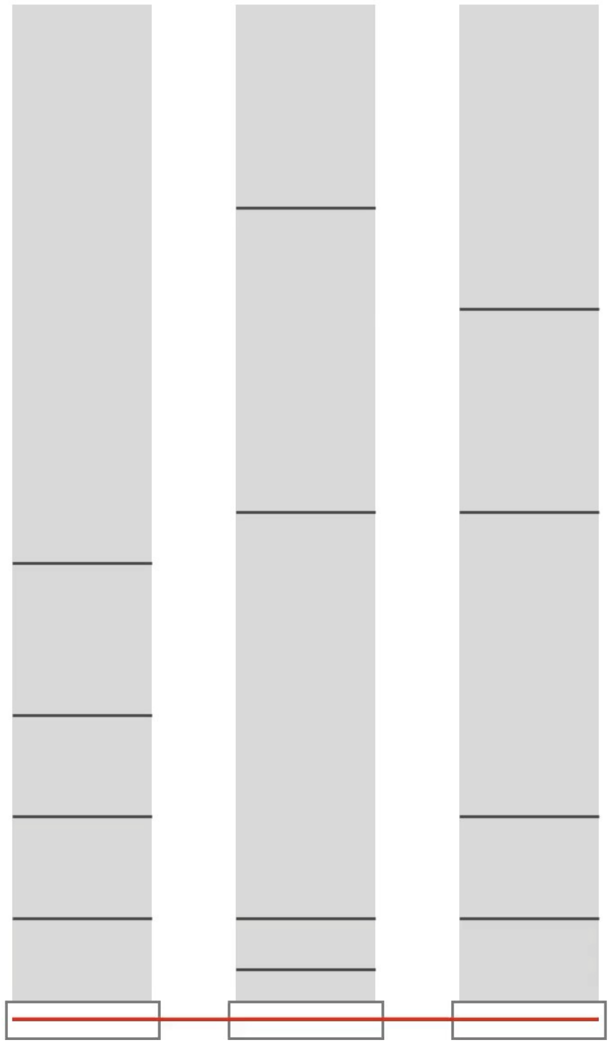

ℓ_1 Mean Approximation for Ladder

- Bounding ℓ_∞ distance also yields ℓ_1 bounds:

  $m = 3 \Rightarrow$ (2/3)-approximate

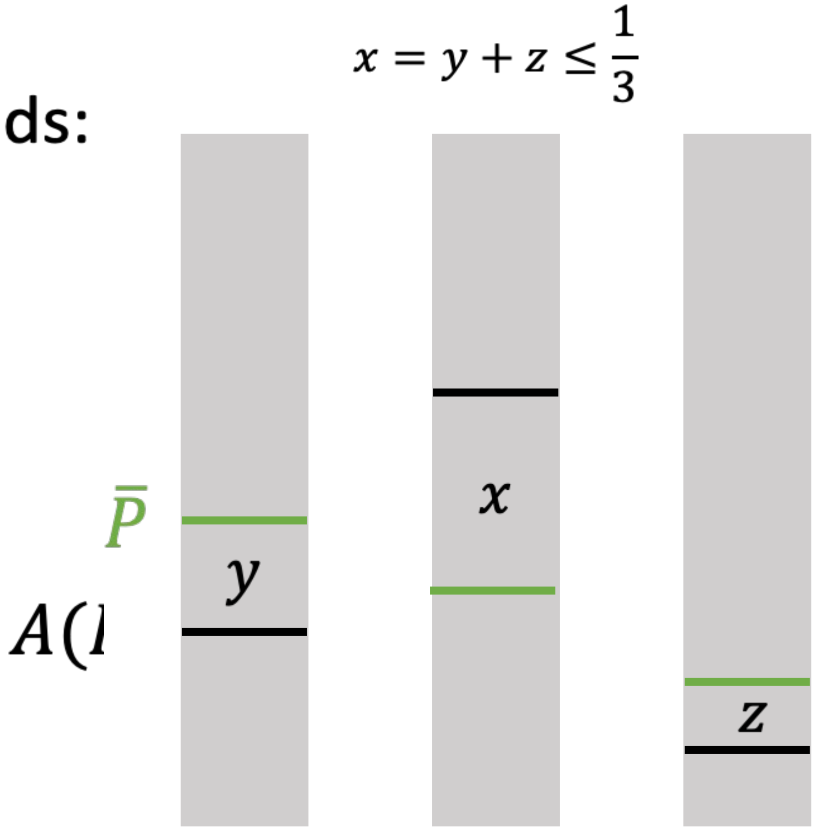

Open question: Can the bounds for $m > 3$ be improved?

Outline

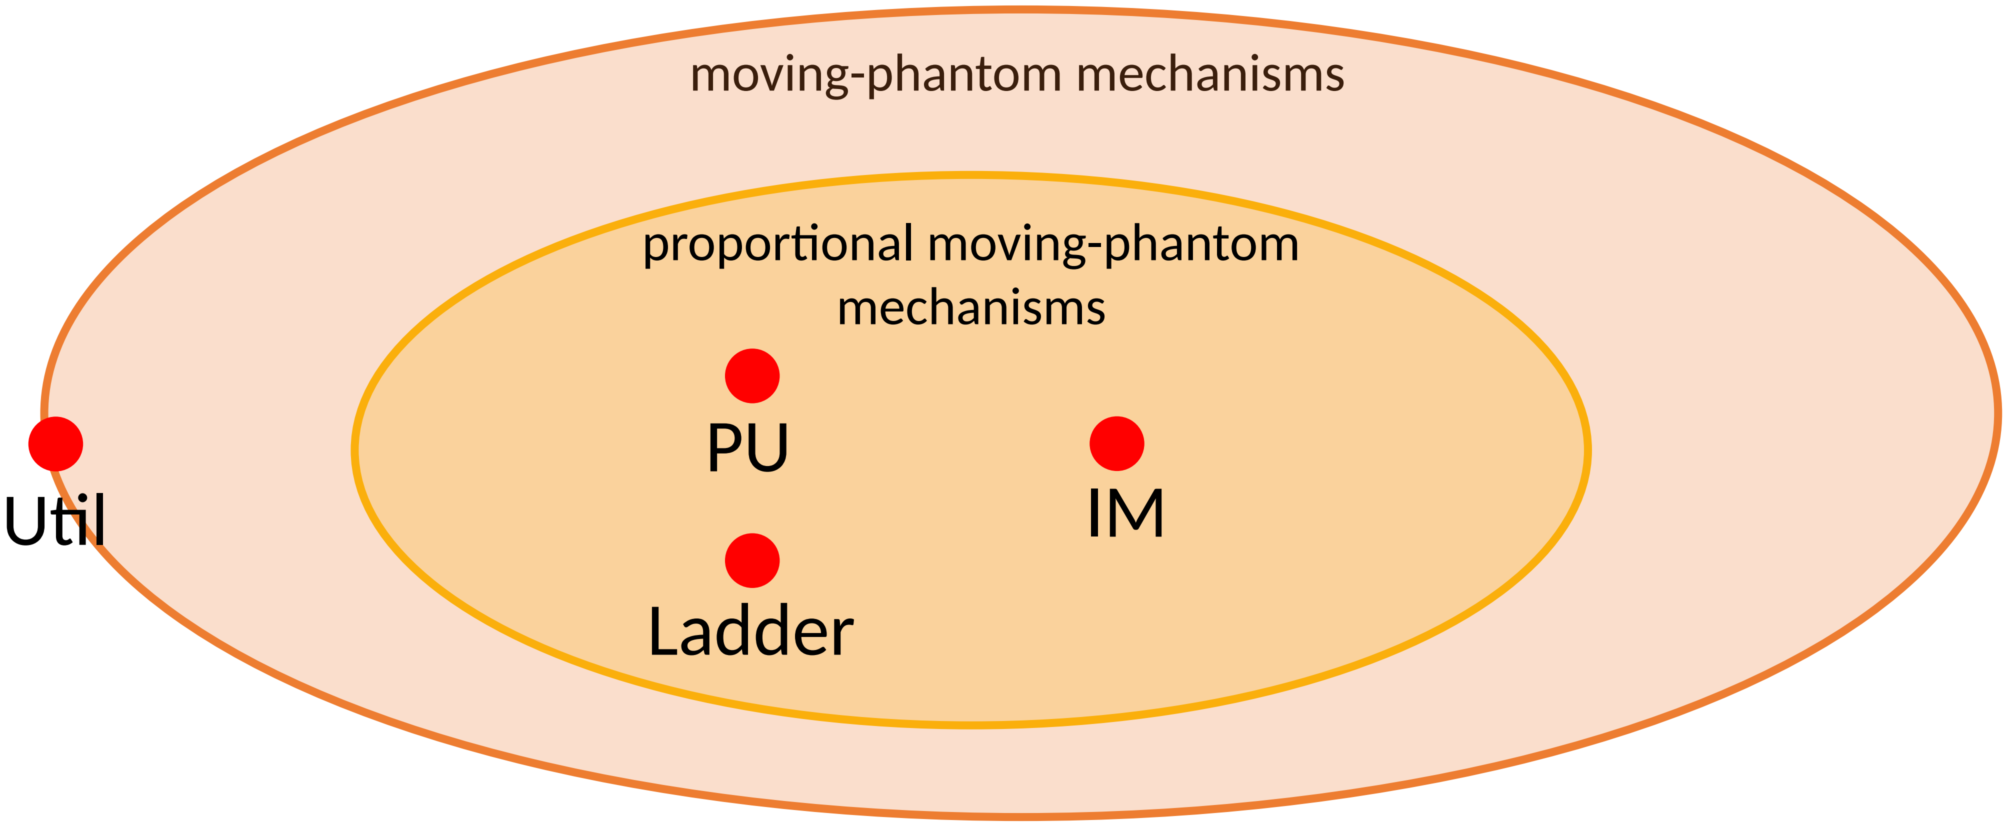
moving-phantom mechanisms
proportional moving-phantom mechanisms
PU
IM
Ladder
Util

“Beyond” PO [Cembrano, F., Schmidt-Kraepelin, Utke 2026]

- Pareto optimality is incompatible with single-minded proportionality. Can we get high social welfare + proportionality?
- UtilProp: Move phantoms one at a time to their proportional positions
  - Same as Util but with phantom positions capped to impose proportionality
- Theorem: UtilProp attains social welfare at least $\frac{2\sqrt{n}-1}{n}$ as high as Util. No single-minded proportional mechanism can do better.

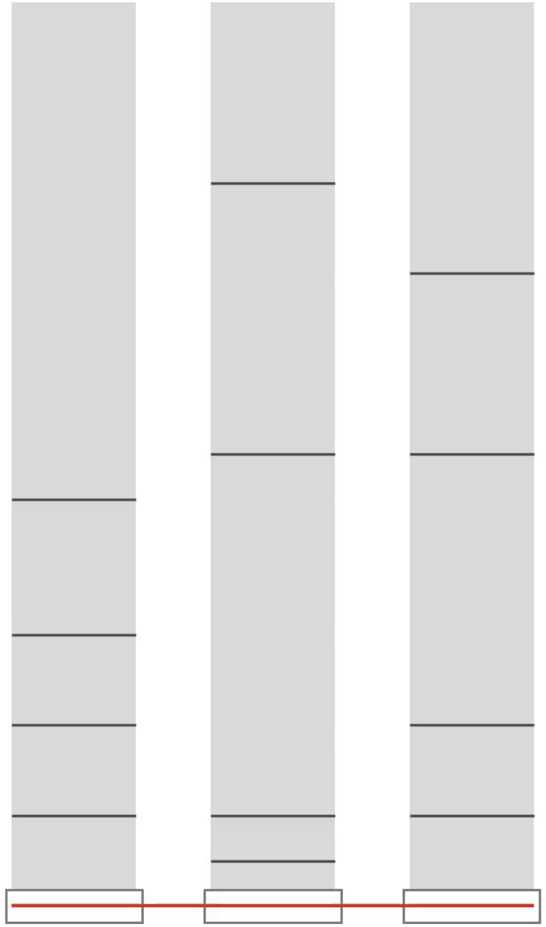

Outline

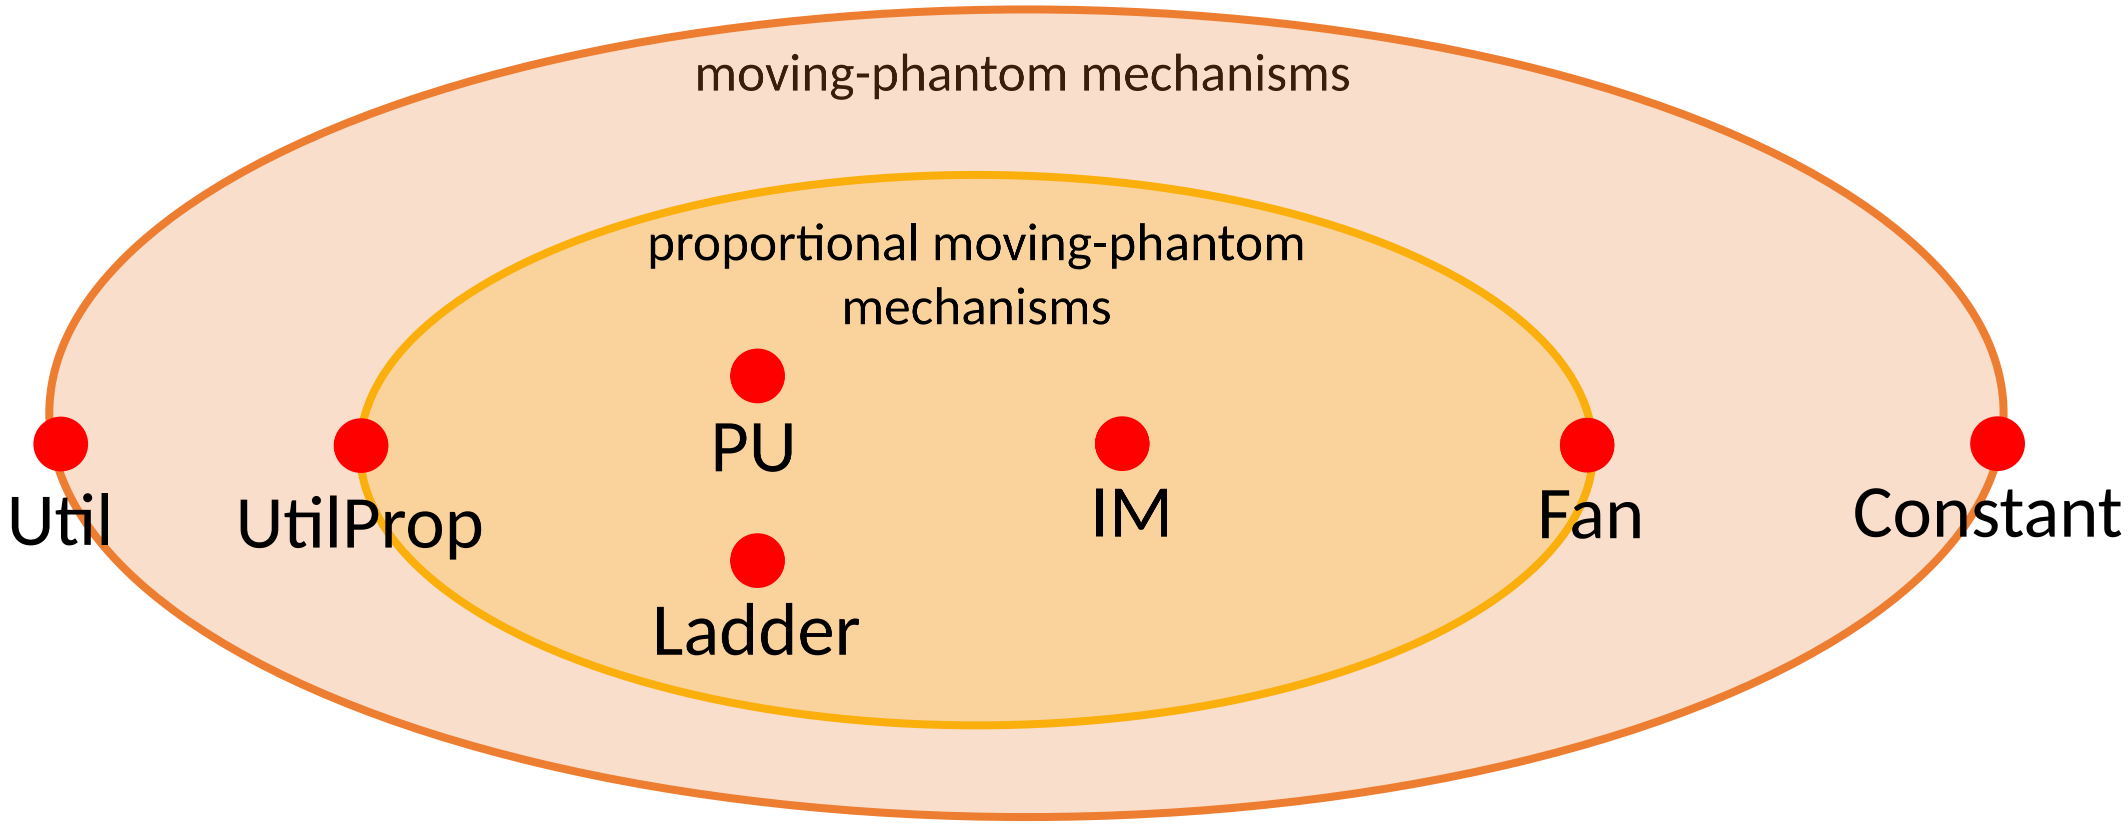

moving-phantom mechanisms
proportional moving-phantom mechanisms
PU
Ladder
Util
UtilProp
IM
Fan
Constant

Per-Instance Welfare Comparisons

- Theorem (informal): If, at time of normalization, phantom positions for mechanism A can be obtained from phantom positions of mechanism B by shifting all phantoms above some threshold up and all phantoms below the threshold down, then A has higher social welfare than B.

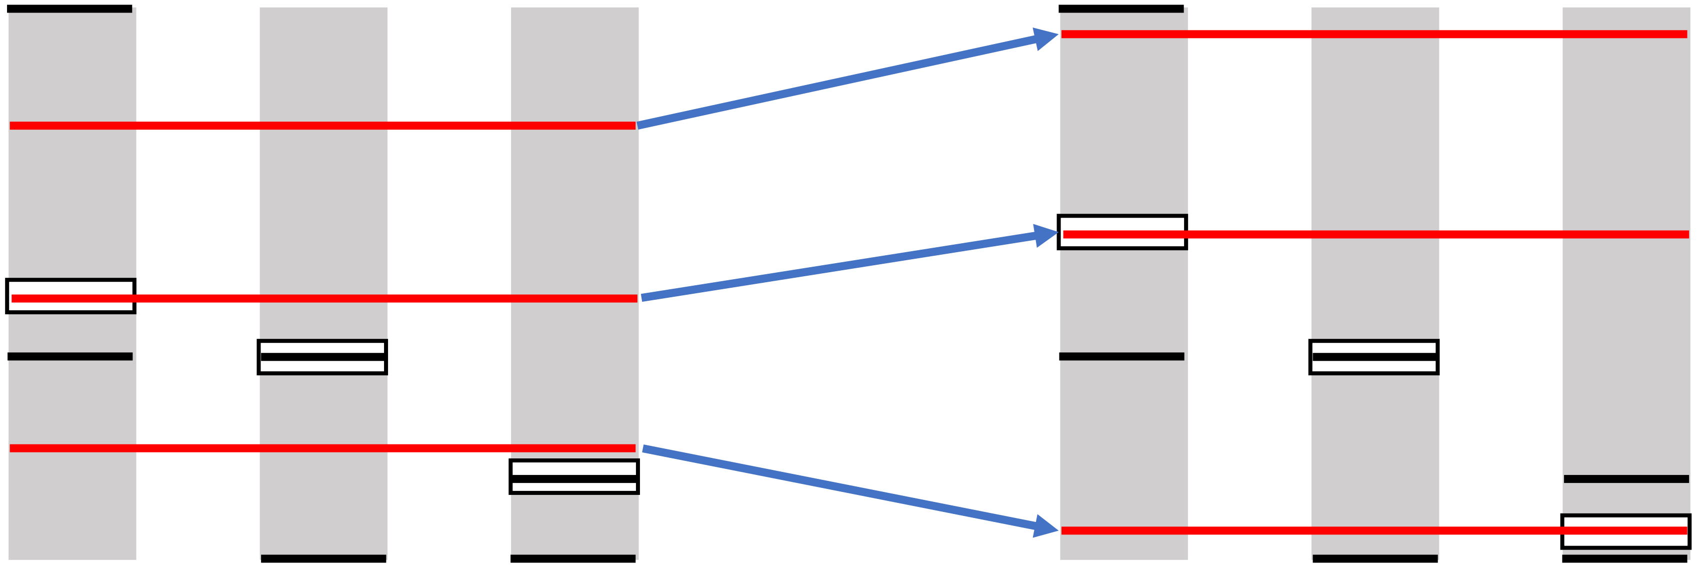

A Partial Ordering over Moving-Phantoms

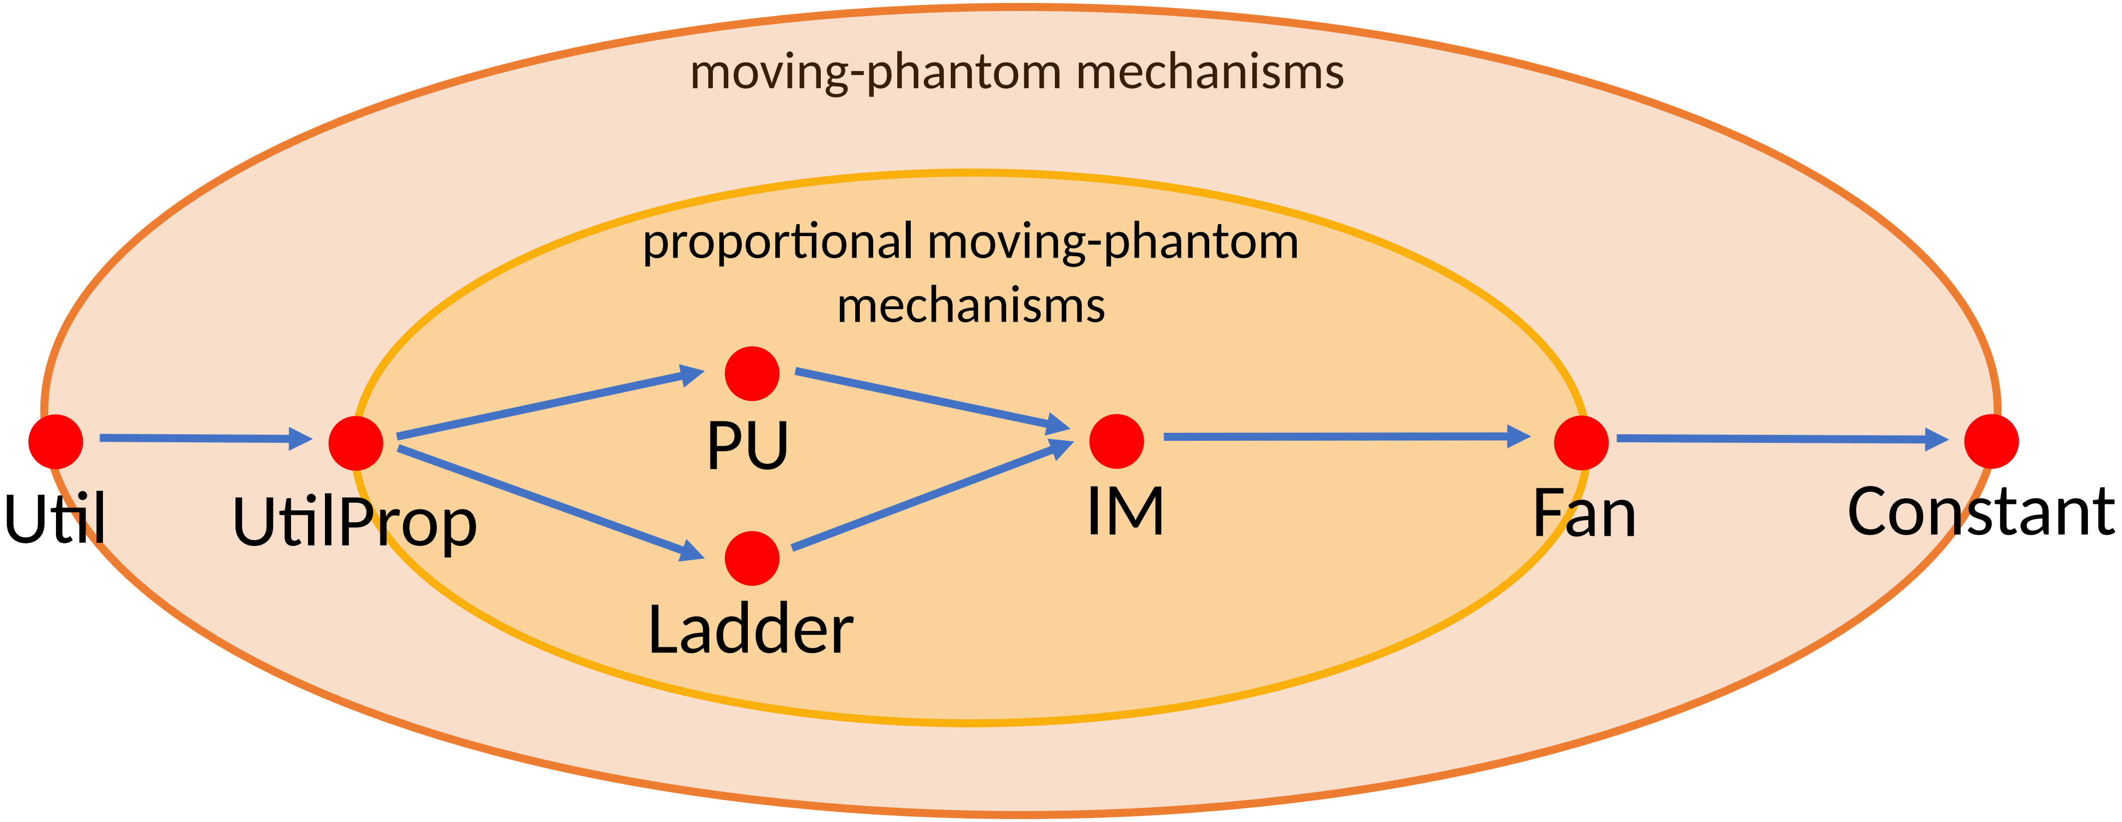

Open question: Does this partial ordering form a lattice?

More mechanisms? [de Berg, F., Schmidt-Kraepelin, Utke 2024]

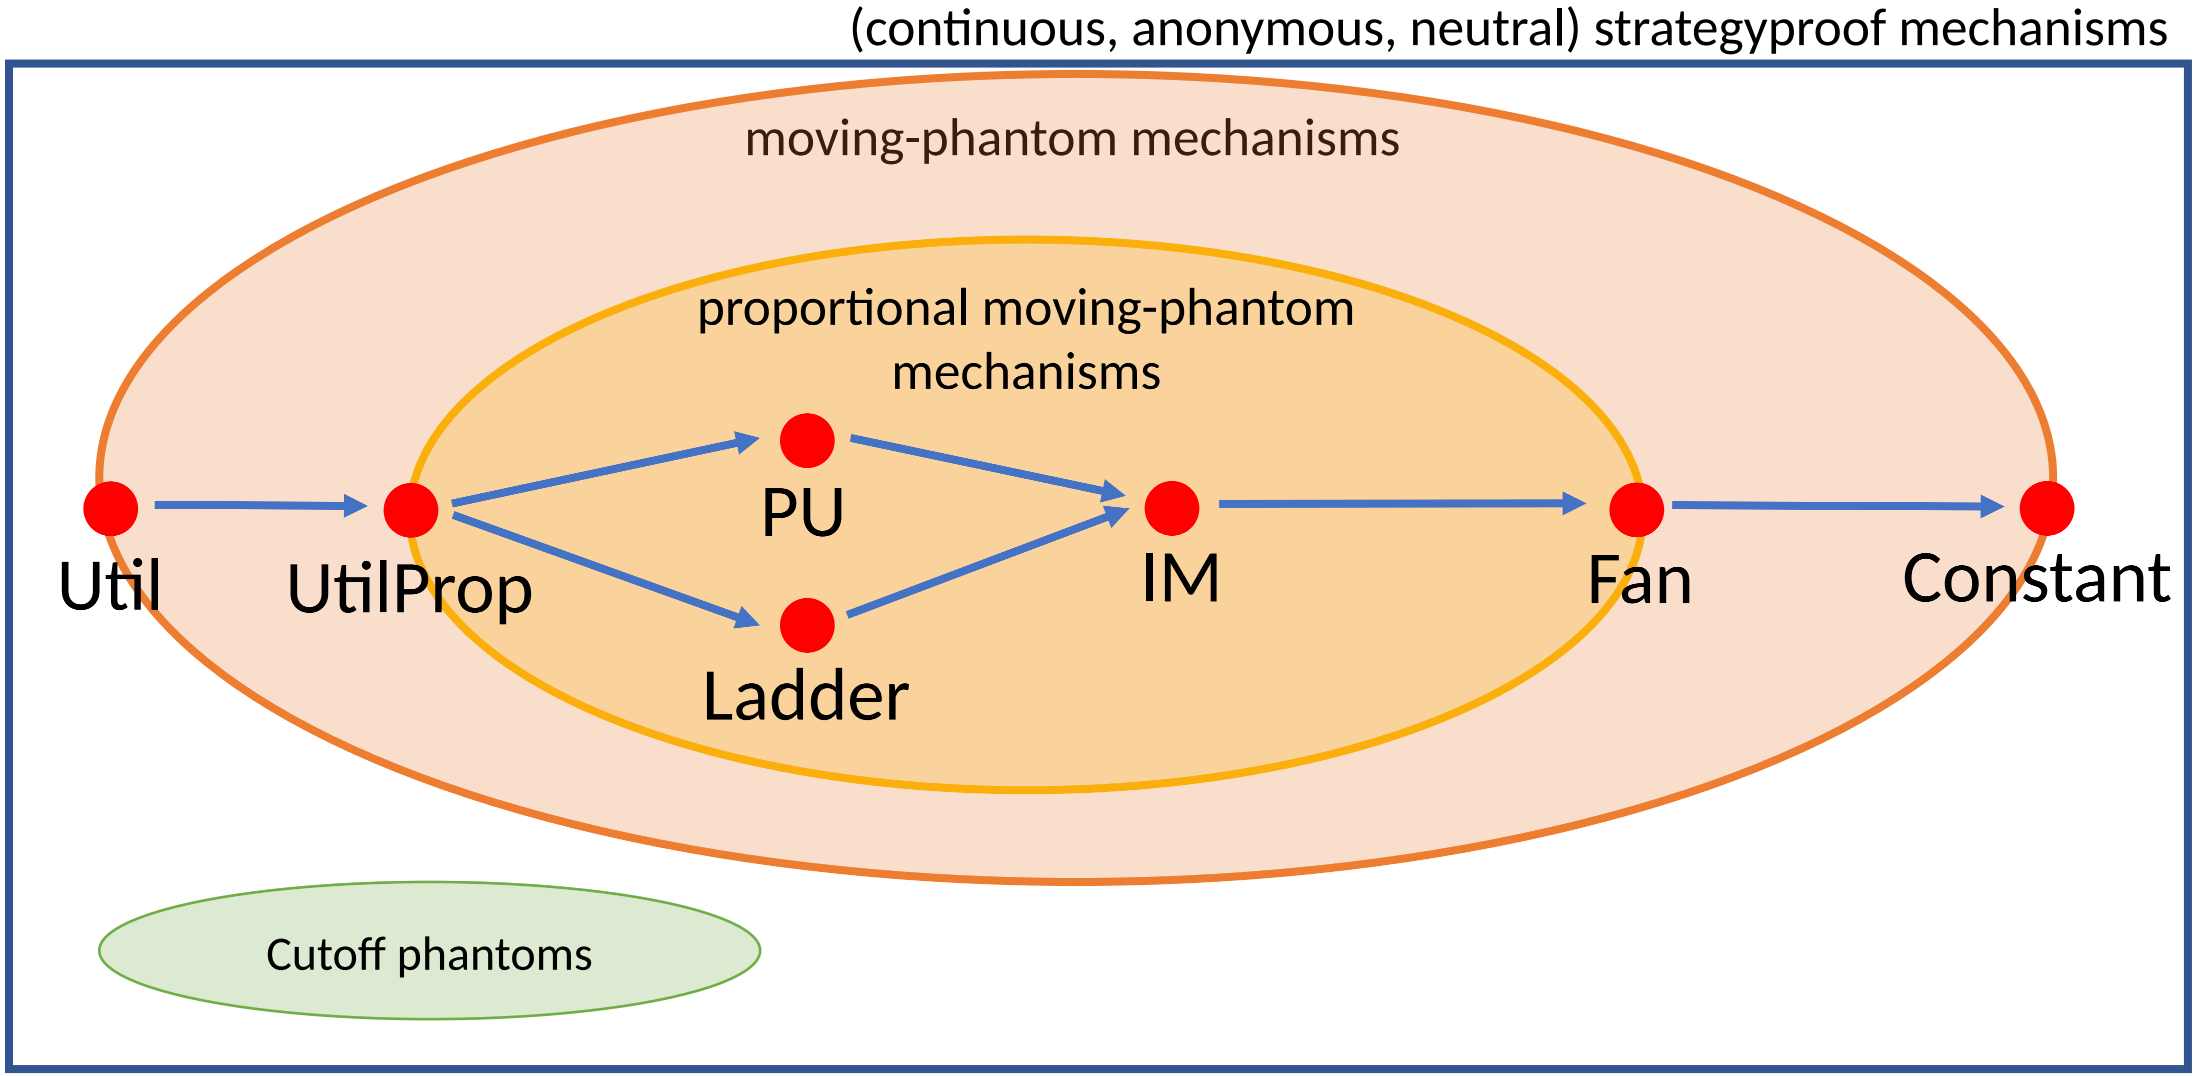

References

[1] Moulin, H. (1980). *On Strategy-Proofness and Single-Peakedness*. Public Choice

[2] Lindner, T.,, Nehring, K., & Puppe, C. *Allocating public goods via the midpoint rule*. (2008). Proc SCW Conference.

[3] Goel, A., Krishnaswamy, A. K., Sakshuwong, S., & Aitamurto, T. (2019). *Knapsack voting for participatory budgeting*. ACM TEAC.

[4] Nehring, K., & Puppe, C. *Resource allocation by frugal majority rule*. (2019). Working Paper.

[5] Freeman, R., Pennock, D. M., Peters, D., & Wortman Vaughan, J. (2021). *Truthful Aggregation of Budget Proposals*. JET.

[6] Caragiannis, I., Christodoulou, G., & Protopapas, N. (2022). *Truthful Aggregation of Budget Proposals with Proportionality Guarantees*. AAAI.

[7] Freeman, R., & Schmidt-Kraepelin, U. (2024). *Project-Fair and Truthful Mechanisms for Budget Aggregation*. AAAI.

[8] de Berg, M., Freeman, R., Schmidt-Kraepelin, U., & Utke, M. (2024). *Truthful Budget Aggregation: Beyond Moving-Phantom Mechanisms*. WINE.

[9] Brandt, F., Greger, M.,. Segal-Halevi, E., & Suksompong, W. (2025). *Optimal Budget Aggregation with Star-Shaped Preference Domains*. MOR.

[10] Cembrano, J., Freeman, R., Schmidt-Kraepelin, U., & Utke, M. (2026). *Social Welfare in Budget Aggregation*. In submission.