Strategyproof Peer Selection

Haris Aziz, Omer Lev, Nicholas Mattei,
Jeffrey S. Rosenschein & Toby Walsh

NSF current state

Description of the Merit Review Process

…

- Selecting **reviewers and panel** members…

- Checking for **conflicts of interest**. In addition to checking proposals and selecting reviewers with no apparent **potential conflicts**, NSF staff members provide reviewers guidance and instruct them how to identify and declare potential conflicts of interest.

NSF proposal

Preliminary Proposals for Core Programs

The mechanism design approach to proposal review is based on the mathematical theory of games, or, more precisely, reverse game theory, namely how the rules of the game should be designed in order to obtain certain desired goals…
the **reviewers assigned from among the set of PIs whose proposals are being reviewed**… Each proposal is assigned for review to seven otherwise non-conflicted PIs … The reviewers must provide both a written review and an **ordering of the seven proposals** to which they are assigned…

NSF proposal

The score of the PI's own proposal is then supplemented with "**bonus points**" depending upon the **degree to which his or her ranking agrees with the consensus ranking**. The award of bonus points is the step that game theory suggests should provide an incentive to each reviewer to give a fair and thorough rating and ranking of the proposals to which he or she is assigned.

NSF problems

Bad reviewers?

Incentive for consensus

Incentive to lower good papers' grade

Laziness

NSF problems

Bad reviewers?

Incentive for consensus

Incentive to lower good papers' grade

Laziness

The model

A set of candidates C={1,…,n}

The model

A set of candidates C={1,…,n}

A set of voters V={1,…,n}

The model

A set of agents $N=\{1,\ldots,n\}$

The model

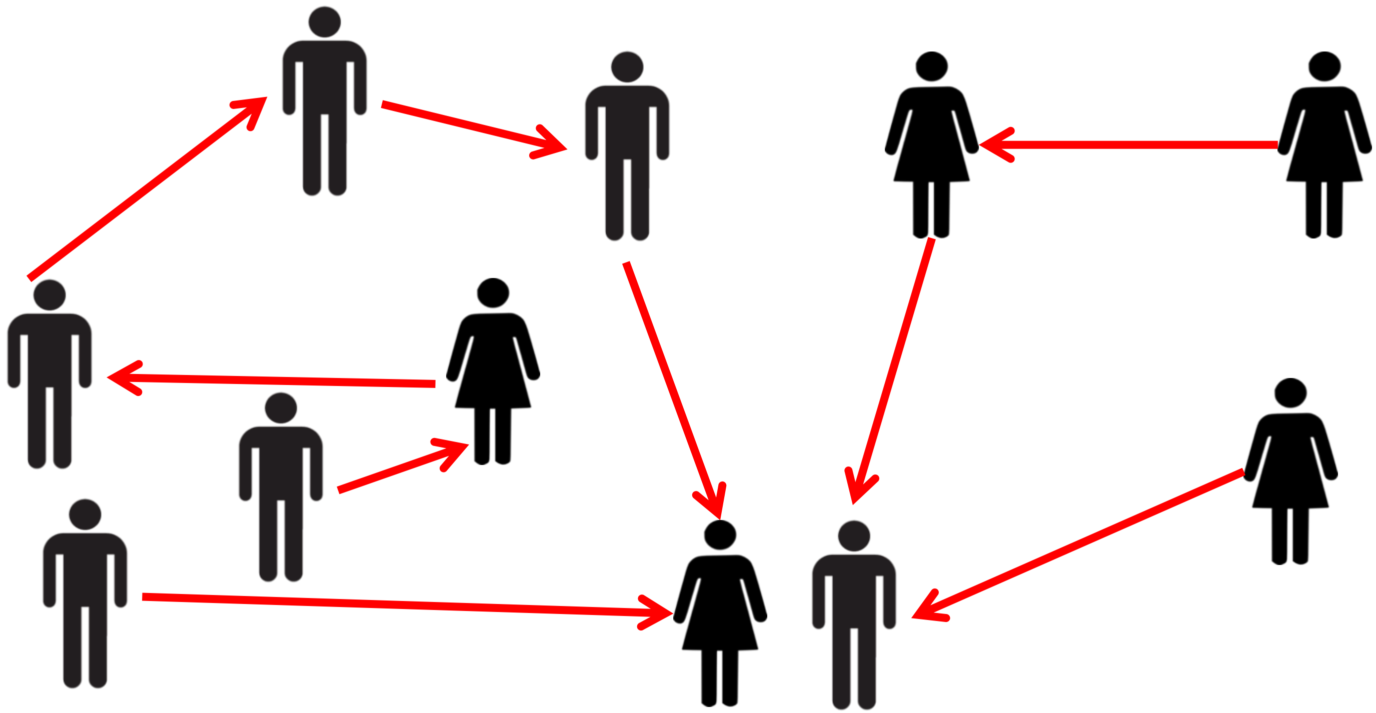

A set of agents N={1,…,n}

Each agent grading/ranking *m* other agents

The model

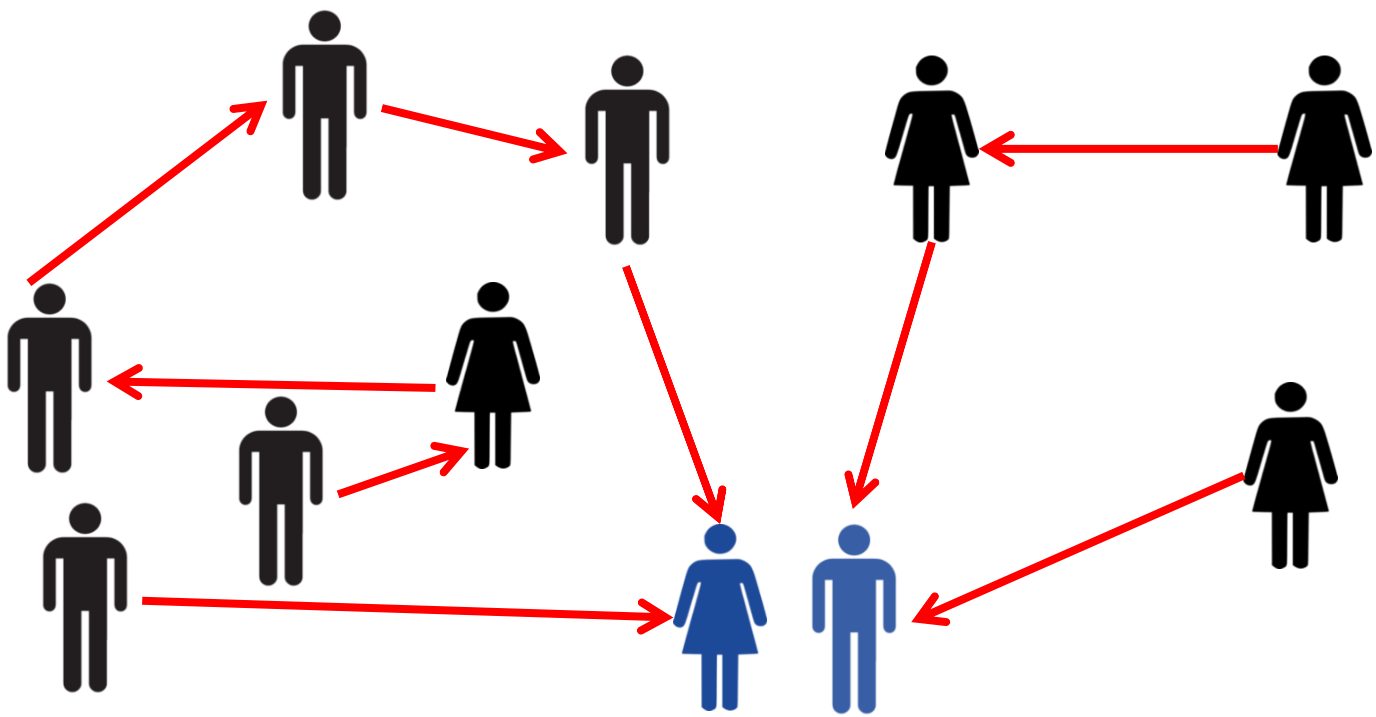

A set of agents N={1,…,n}

We want to select the top *k* agents

Vanilla mechanism & guarantees

Choose the top scoring *k* agents.

Not strategyproof…

Partition

(Alon, Fischer, Procaccia, Tennenholtz; TARK 2011 and others)

Partition basic idea

Achieving strategyproofness by dividing agents into groups, letting no agents in the same partition rate each other. Each partition is considered independently of the rest.

Partition algorithm

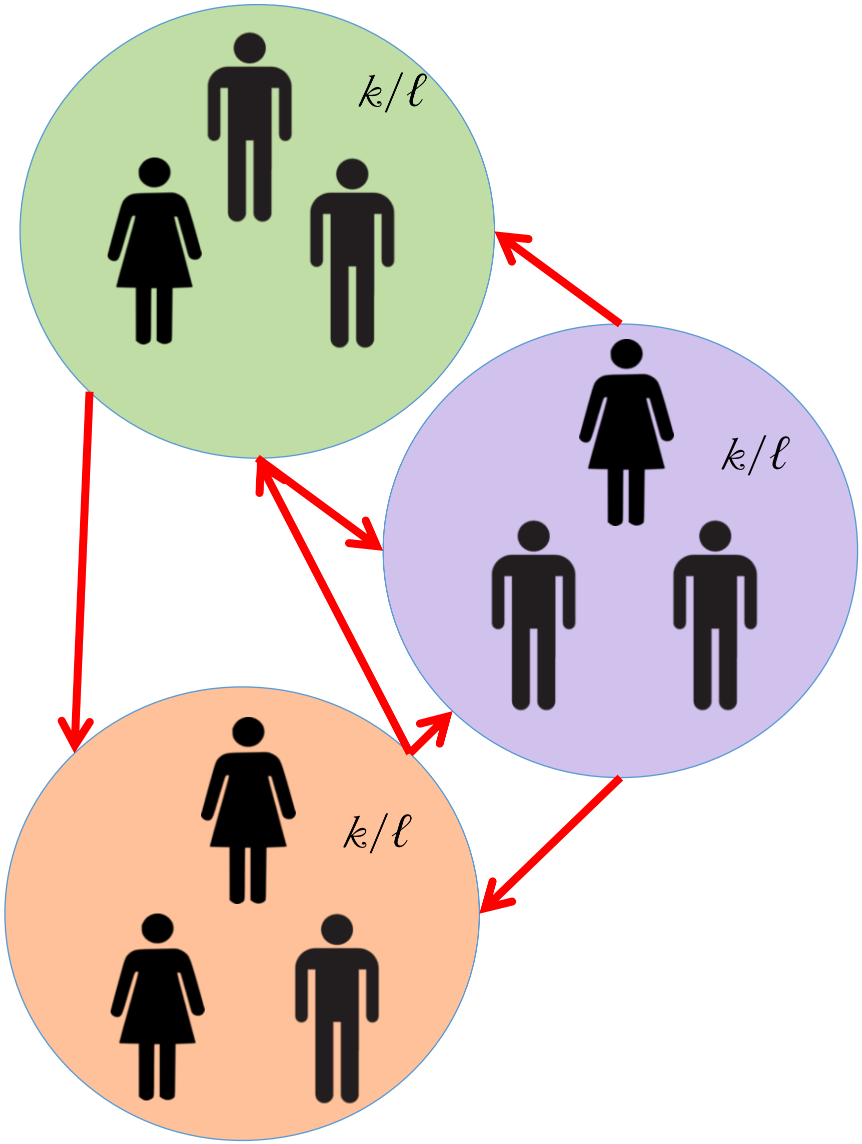

Divide agents to ℓ partitions. Each agent ranks m agents outside their own partition.
Finally, selected agents are the top ranked k/ℓ in each partition.

Why not partition?

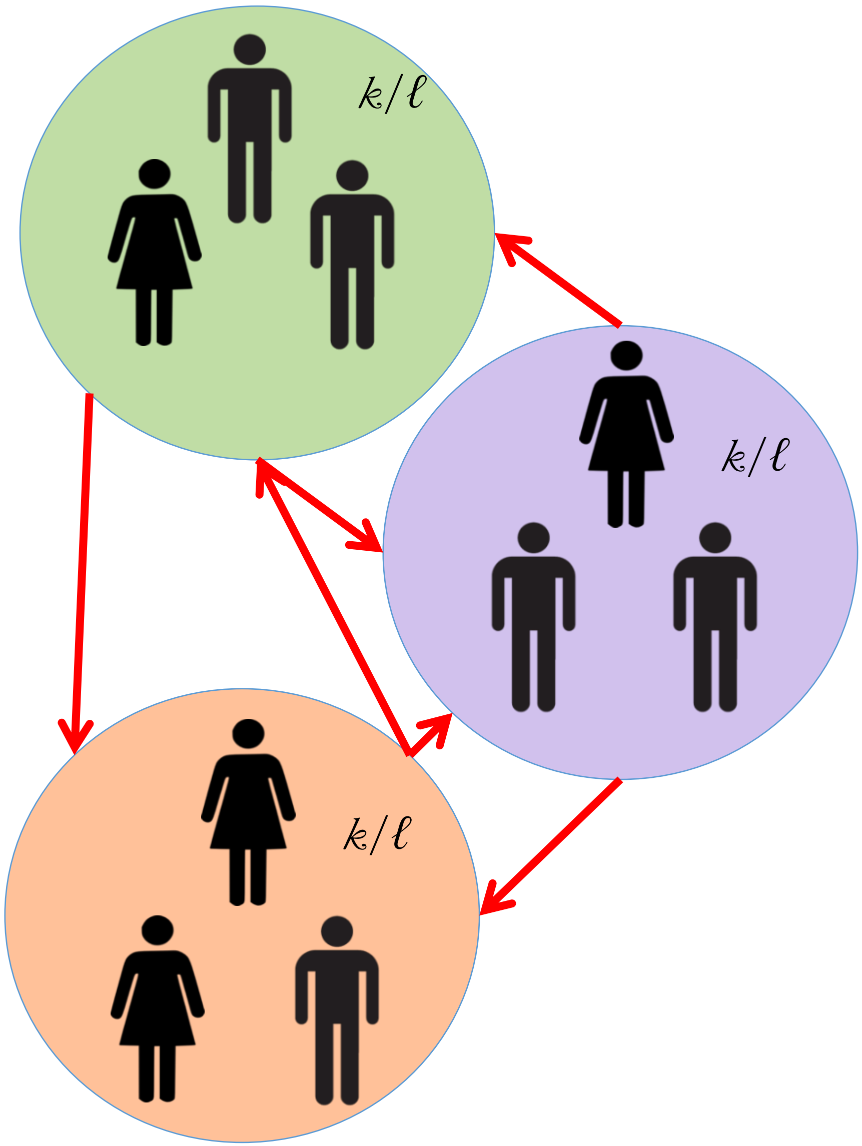

What if one cluster has many good agents, and another has less? Must we treat them equally?

Why not partition?

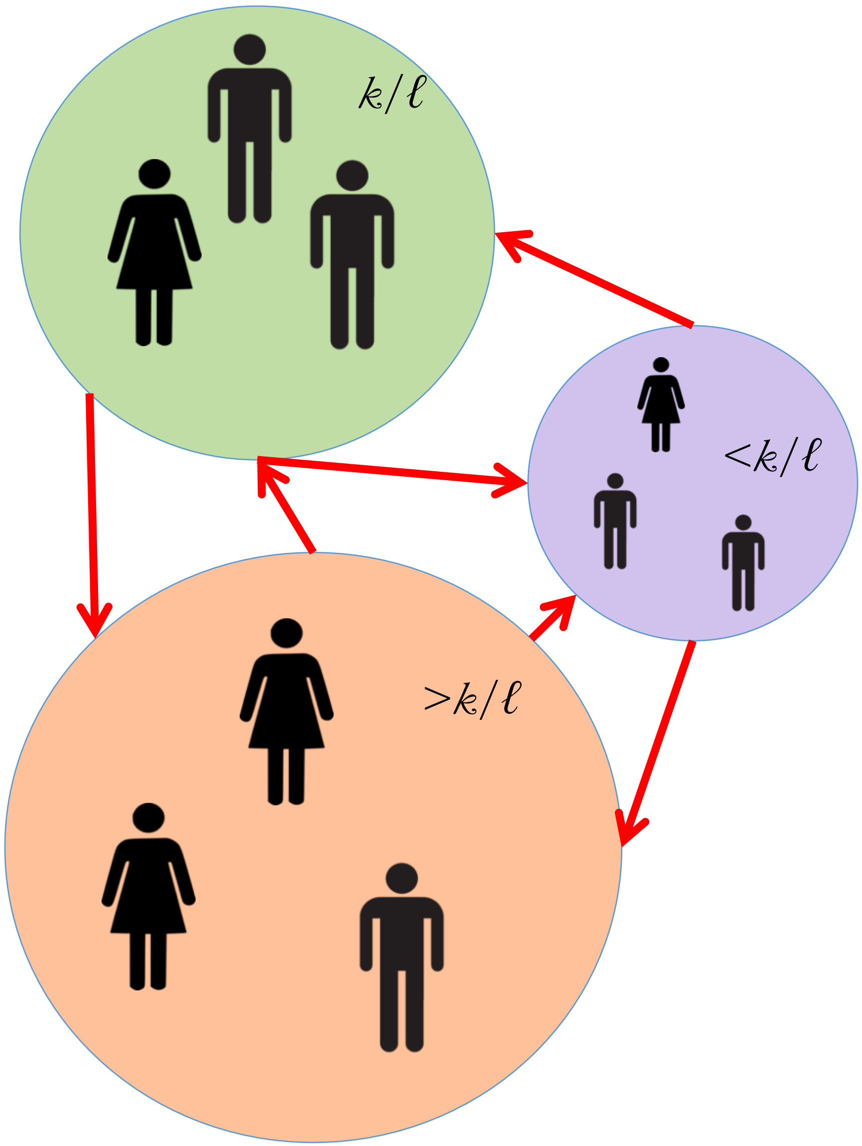

What if one cluster has many good agents, and another has less? Must we treat them equally?

We would like to give them different shares!

Dollar partition

(Aziz, Lev, Mattei, Rosenschein, Walsh; AAAI 2016)

Dollar partition basic idea

Achieving strategyproofness by dividing agents into groups, letting no agents in the same partition rate each other. Each partition ultimate share influenced by its relative strength compared to others.

A small digression…

Dividing a dollar

(de Clippel, Moulin, Tideman; Journal of Economic Theory 2008)

Dividing a dollar problem

Divide a divisible item between agents in a strategyproof manner.

E.g., bonus between employees, based on merit.

Dividing a dollar algorithm

Let each agent divide the dollar between their peers, so for agent i, $\sum_{j \neq i} v_i(j) = 1$.

Ultimately, agent i's share will be

$$x_i = \frac{1}{n} \sum_{j \neq i} v_j(i)$$

Dollar raffle peer selection solution?

Have each agent's share be the probability of it being selected.

Not strategyproof for $k>1$!

Back to our problem…

Dollar partition

Dollar partition algorithm

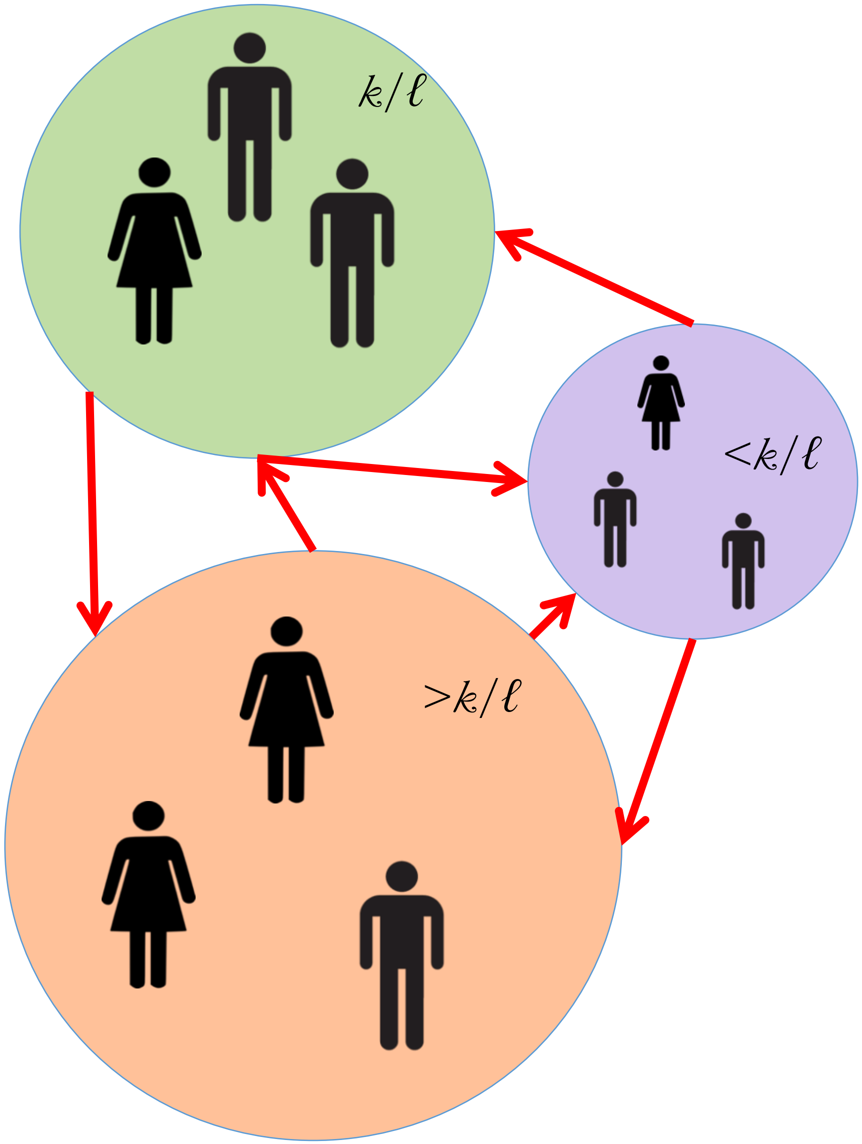

Each agent grades m agents outside their cluster, and we normalize the grades: $\sum_{j \in N} v_i(j) = 1$

Each cluster has a share:

$$x_i = \frac{1}{n} \sum_{j \in C_i, j' \notin C_i} v_{j'}(j)$$

Dollar partition raffle peer selection solution?

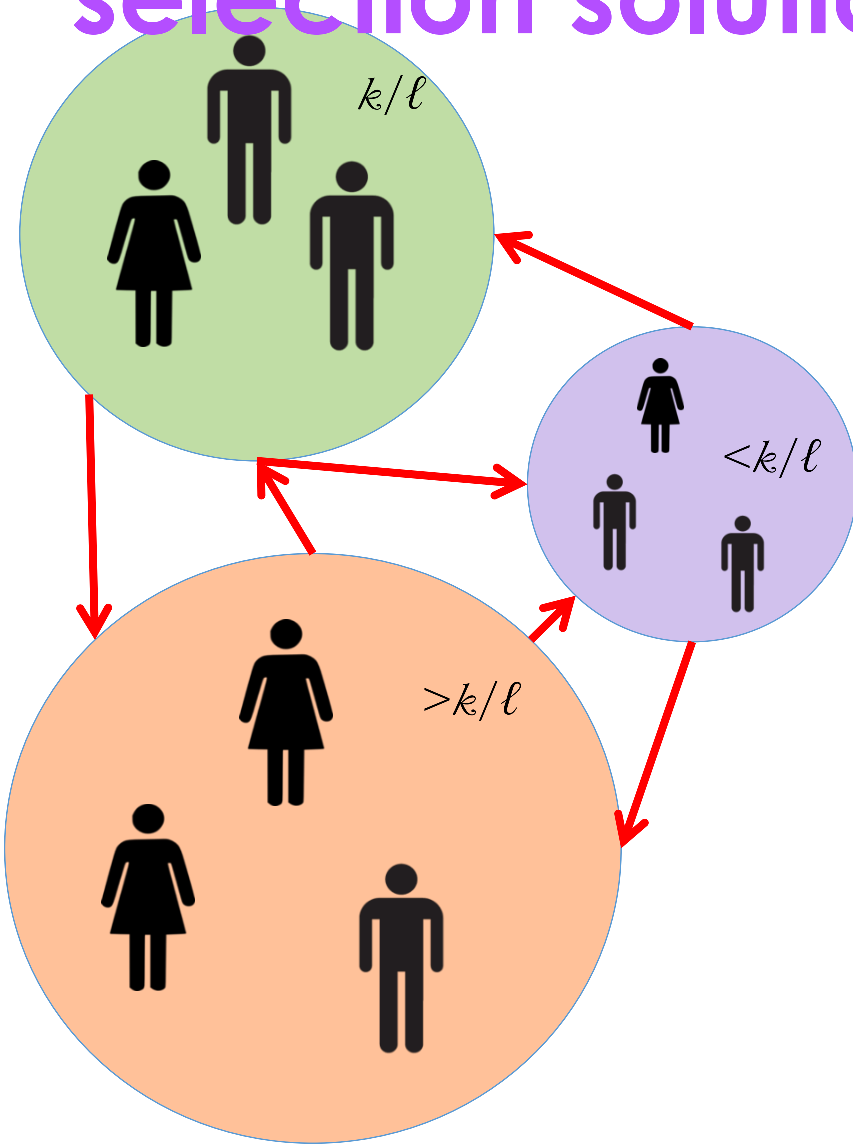

Use shares as probabilities of selecting agents from a cluster?

Could end up selecting all agents from a single cluster…

Dollar partition algorithm

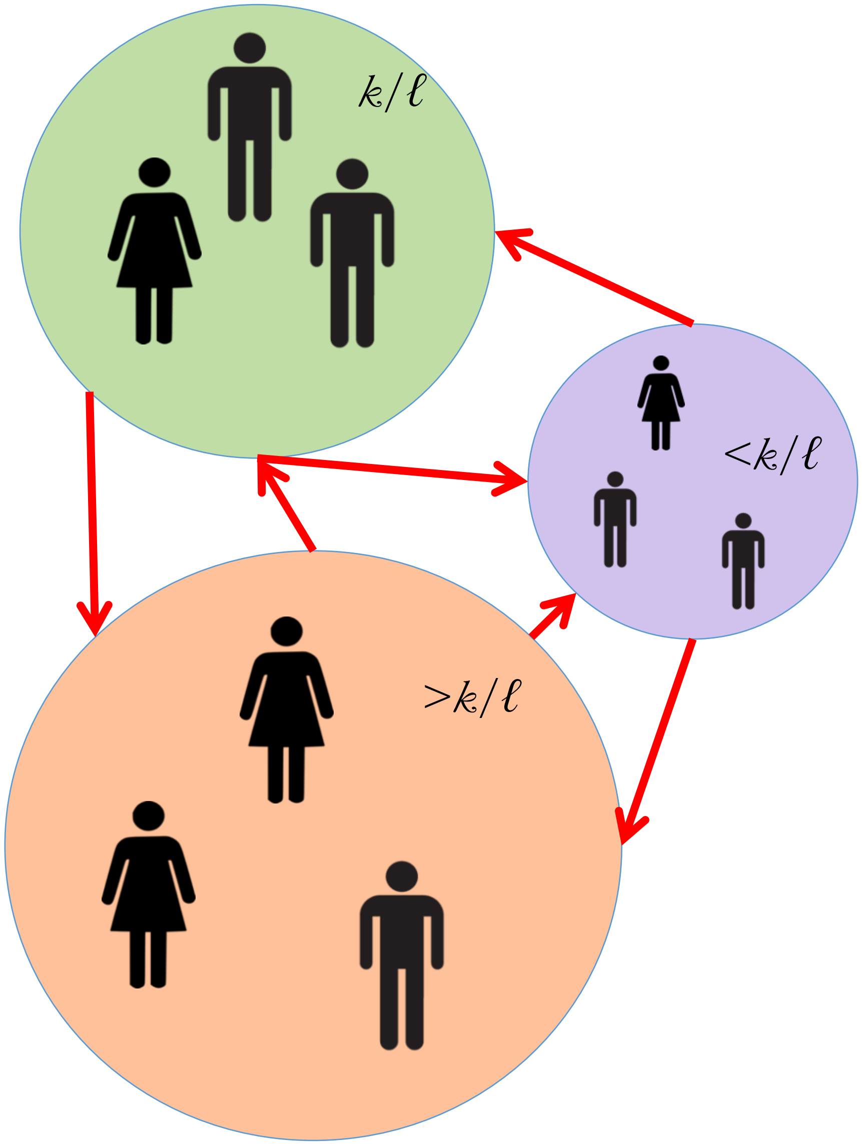

Select the top $k \cdot x_i$ agents from each cluster.

Dollar partition problem

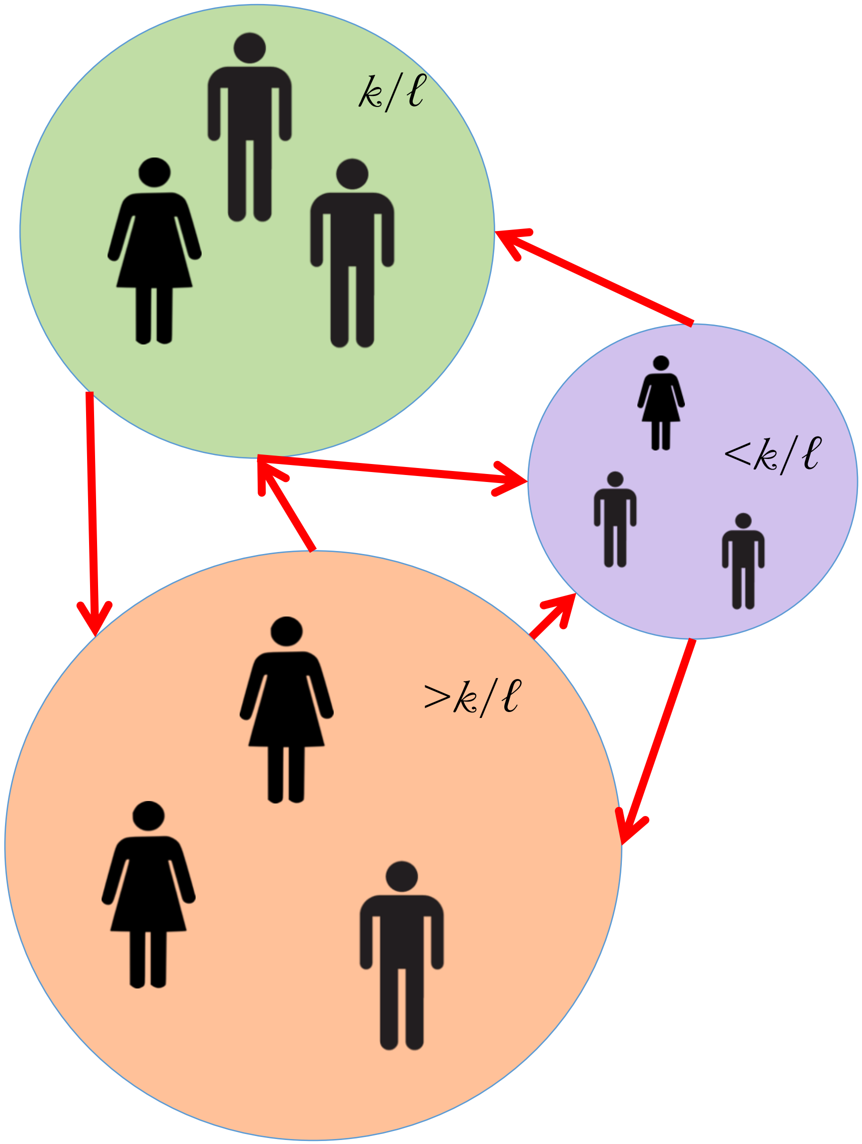

Select the top $k \cdot x_i$ agents from each cluster.

What if $k \cdot x_i$ is a fraction?

Bringing us to…

The allocation problem

Example US ~1790

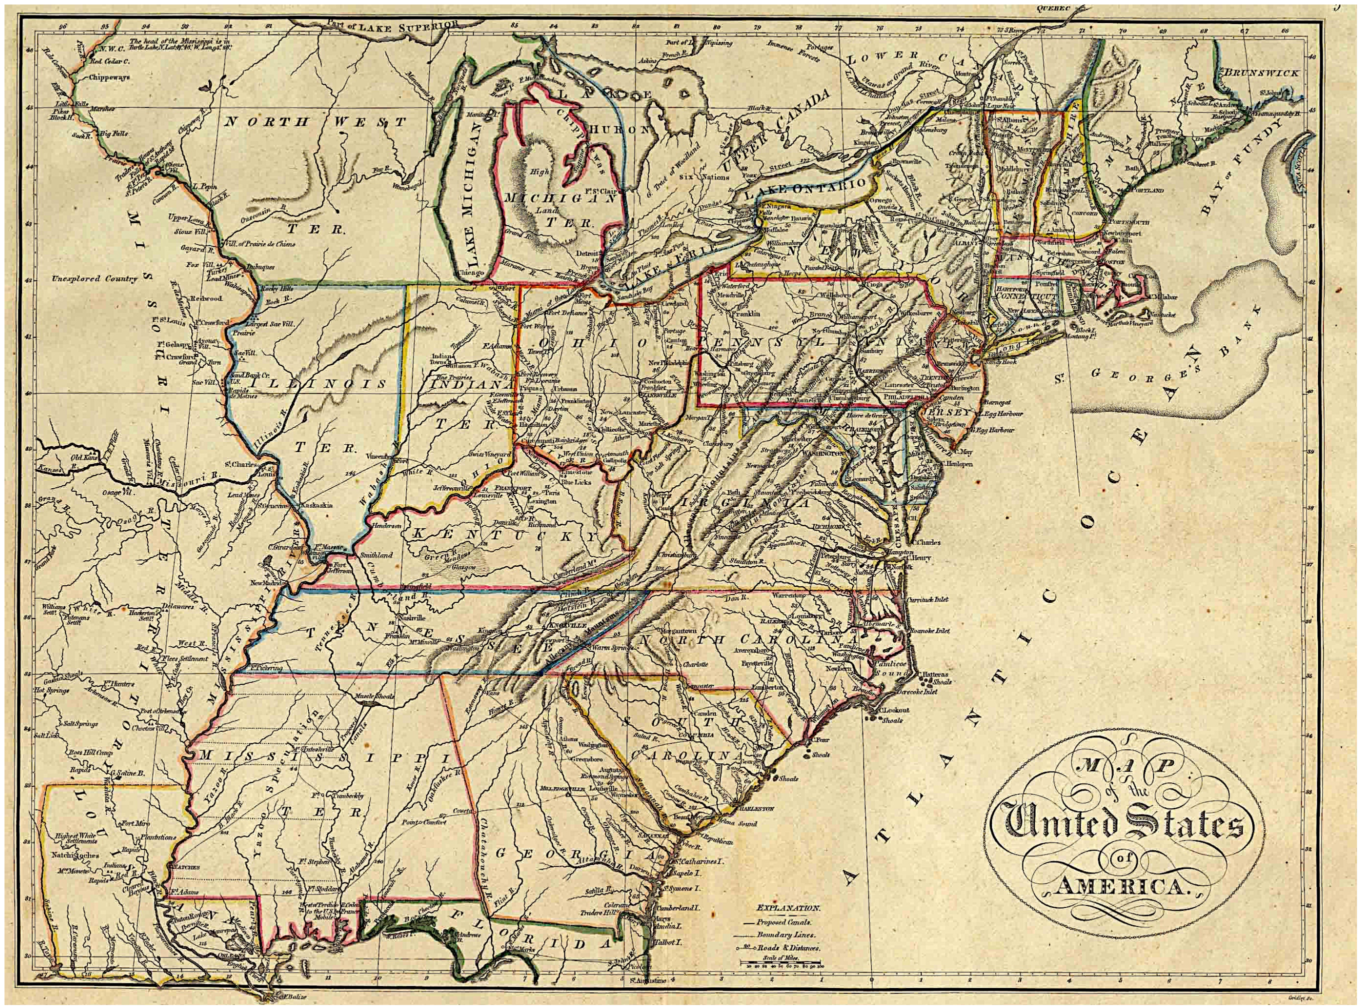

Example US ~1790

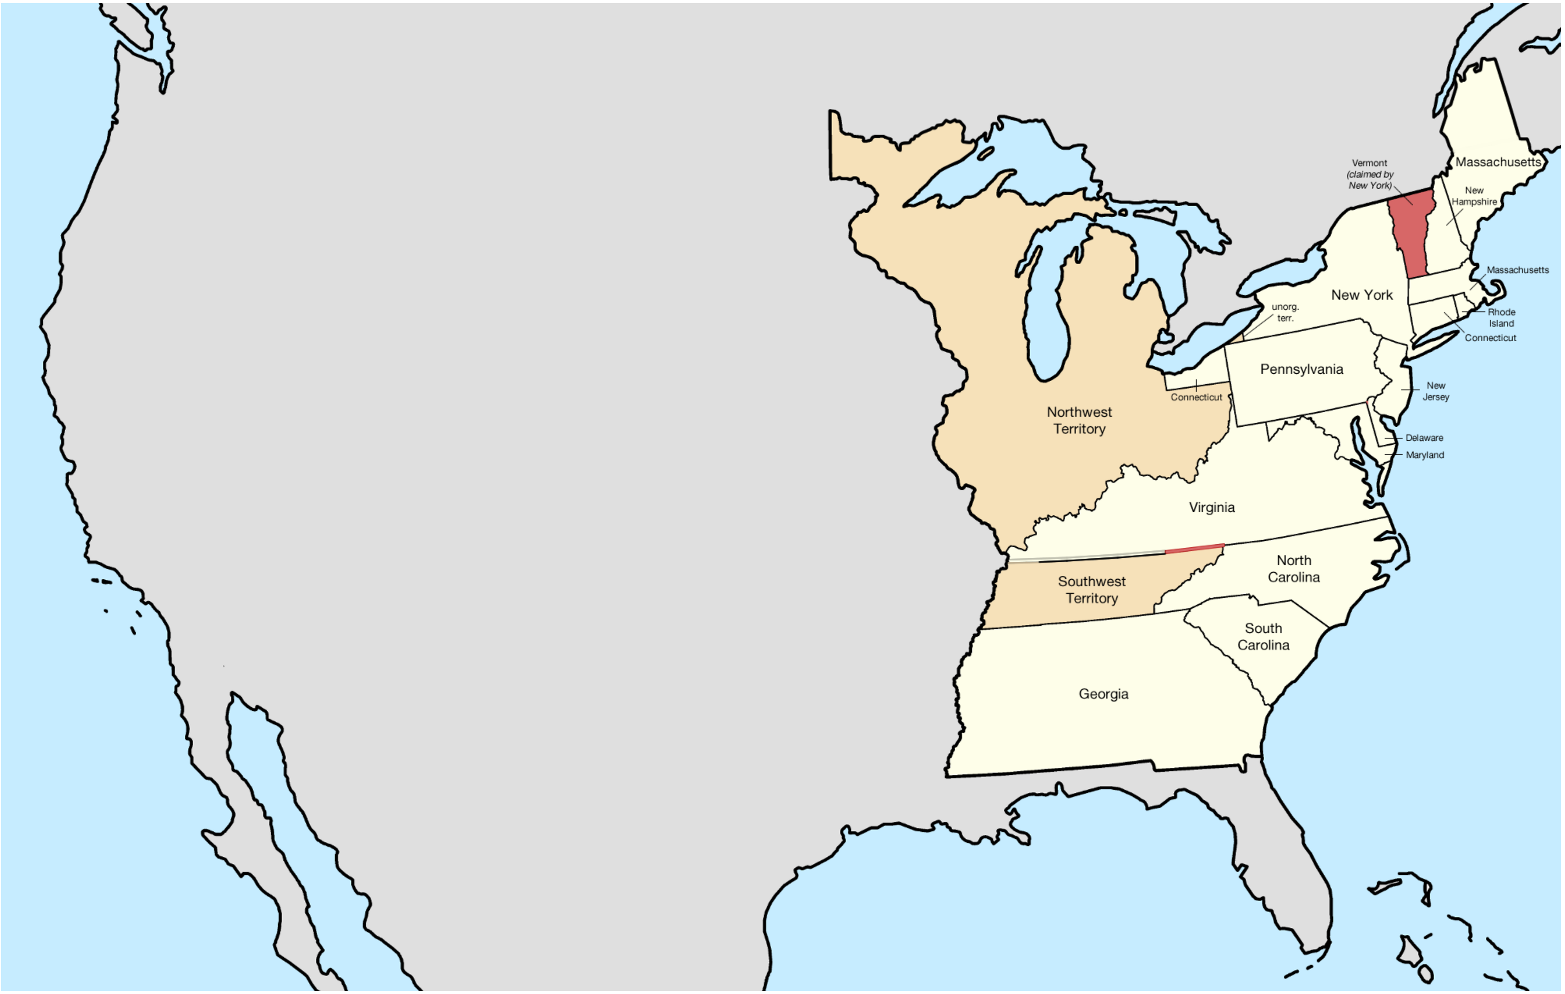

Example US constitution

Article I, section 2:

Representatives and direct Taxes shall be **apportioned among the several States** which may be included within this Union, **according to their respective Numbers**…
The actual Enumeration shall be made within three Years after the first Meeting of the Congress of the United States, and within every subsequent Term of ten Years, in such Manner as they shall by Law direct.

The allocation problem

How to allocate k slots between ℓ clusters, when each cluster has a fractional weight (summing up to k)?

Exact dollar partition

(Aziz, Lev, Mattei, Rosenschein, Walsh; To be submitted…)

Exact dollar partition idea

Achieving strategyproofness by finding an allocation mechanism on top of dollar partition, that lets us select exactly k agents.

Dollar partition algorithm

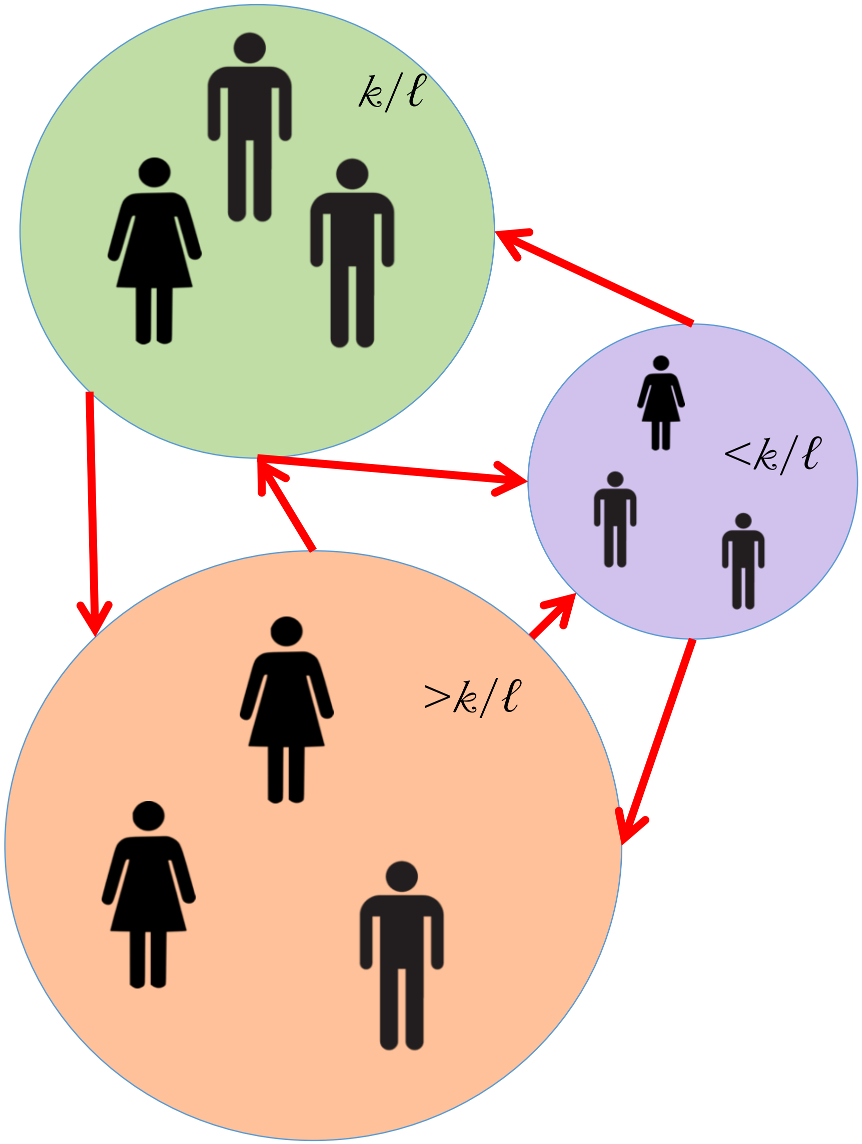

Each agent grades m agents outside their cluster, and we normalize the grades: $\sum_{j\in N} v_i(j) = 1$

Each cluster has a quota:

$$k \cdot x_i = k \cdot \frac{1}{n} \sum_{j\in C_i, j'\notin C_i} v_{j'}(j)$$

The allocation problem theorem

No deterministic method of rounding the quotas that guarantees selection of exactly *k* agents can be strategyproof.

Exact dollar partition allocation mechanism

$k=7$

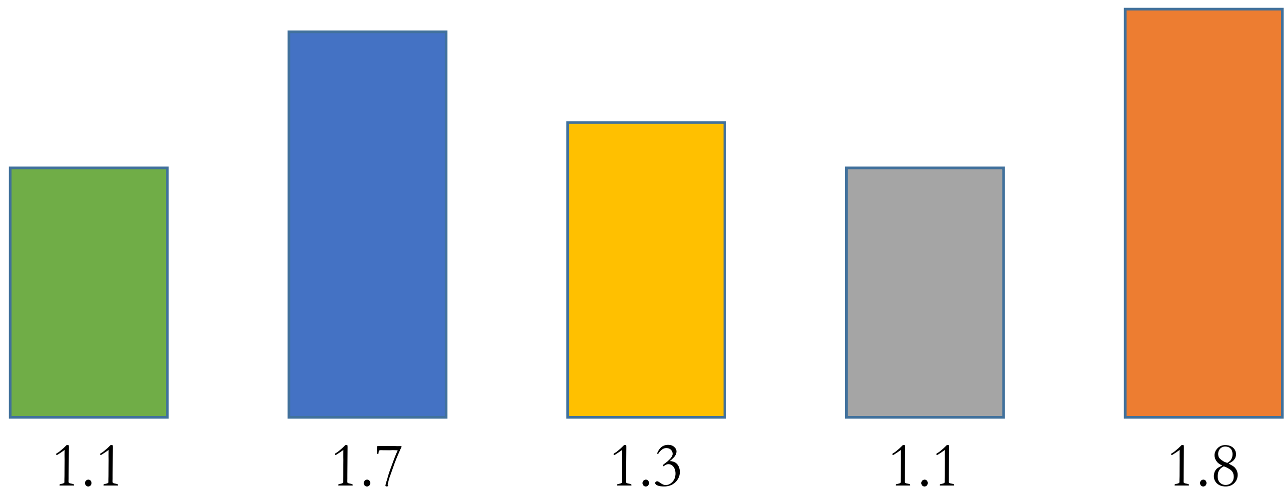

Exact dollar partition allocation mechanism

$k=7$

	2	2	1	1	1	0.1
	1	1	2	2	1	0.1
	1	1	2	1	2	0.2
	1	1	1	2	2	0.6
Expected value:	1.1	1.1	1.3	1.7	1.8	

But which one is best?

(it's exact dollar partition)

Voter preferences
Mallows model

A Mallows model assume the existence of a ground truth, and each agents has a “noisy” version of this truth.

It uses a parameter Φ to indicate distance from the ground truth, indicating the likelihood of a flip from the ground truth.
$\Phi = 0$ means all agents have the ground truth,
$\Phi = 1$ means all agents have randomly assigned preferences.

Voter preferences simulation

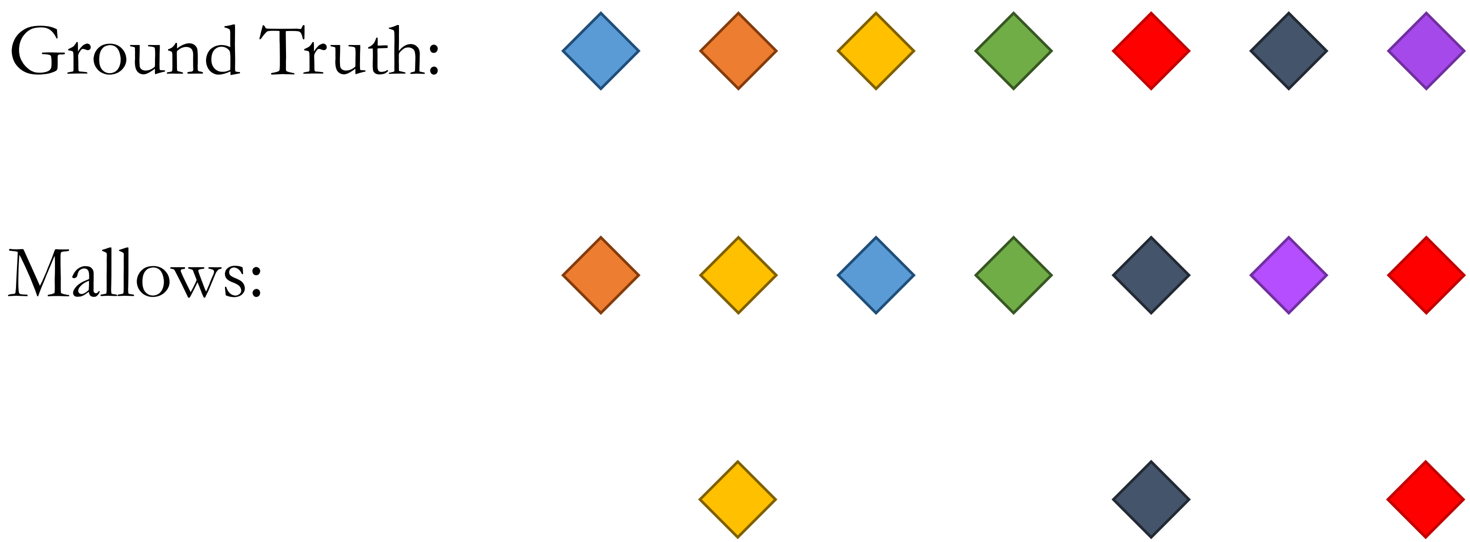

Each agent delivers a partial, noisy preference order.

Setting simulation

Similar setting to the NSF ones, with expanding the parameters.

n: *130* proposals (agents).
m: *5, 7, 9, 11, 13, 15*
ℓ: *3, 4, 5, 6* clusters.
k: *15, 20, 25, 30, 35* winners.
Φ: *0.0, 0.1, 0.2, 0.35, 0.5*
Borda scoring of grades.

Results

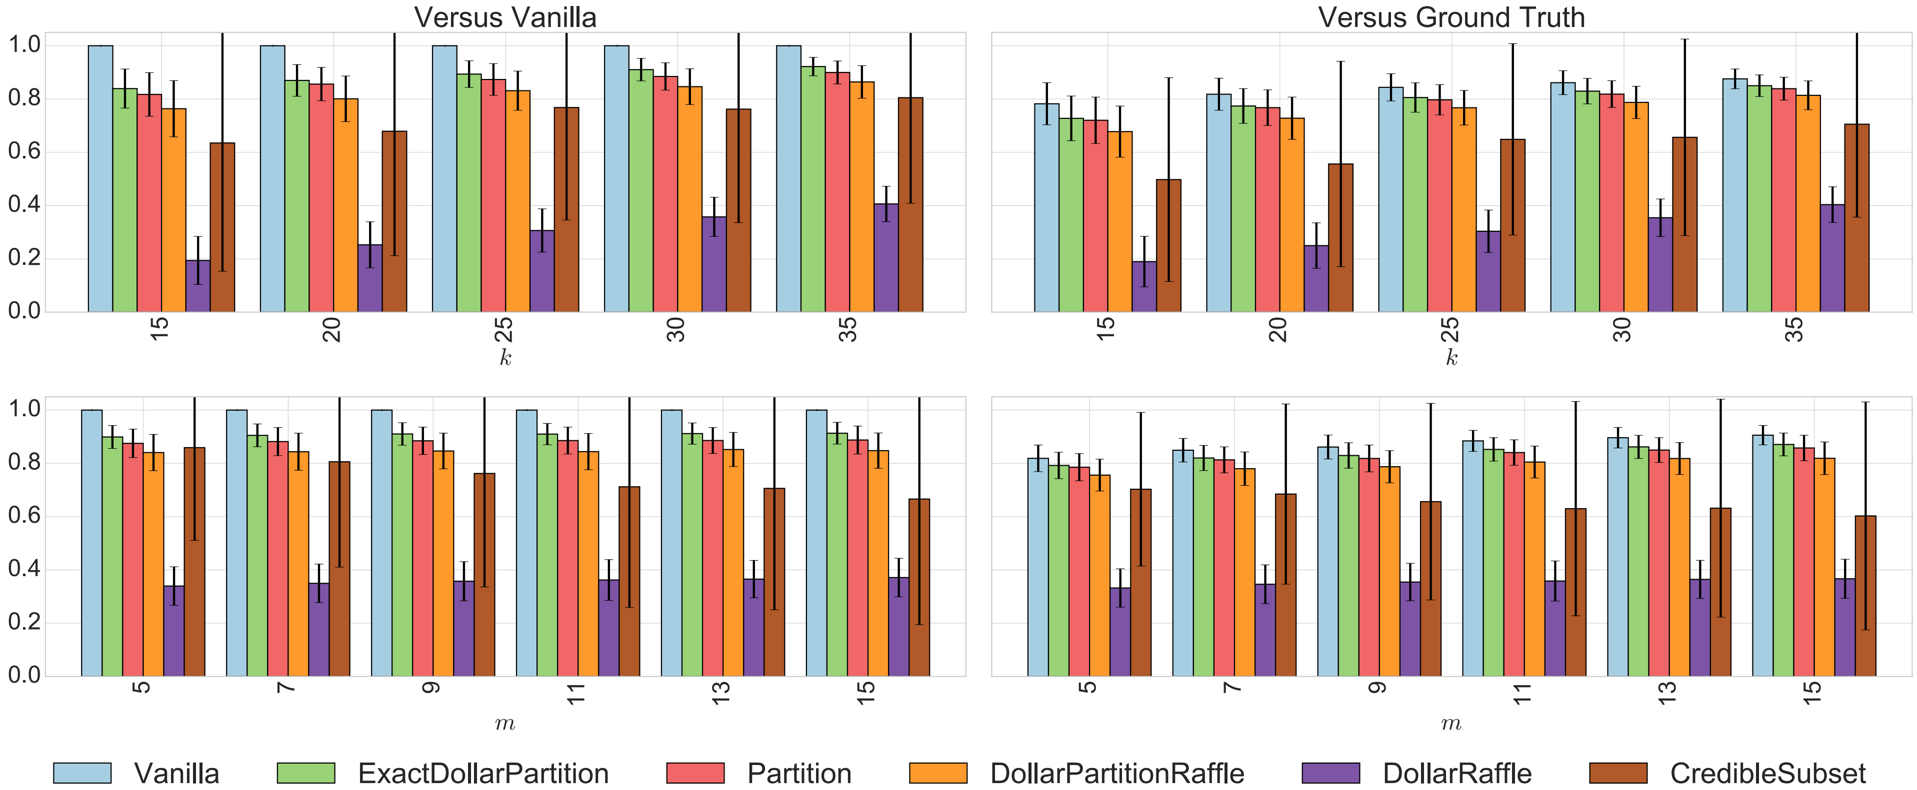

Exact dollar partition **better** than all other Dollar mechanisms and credible subset.

Results vs. partition

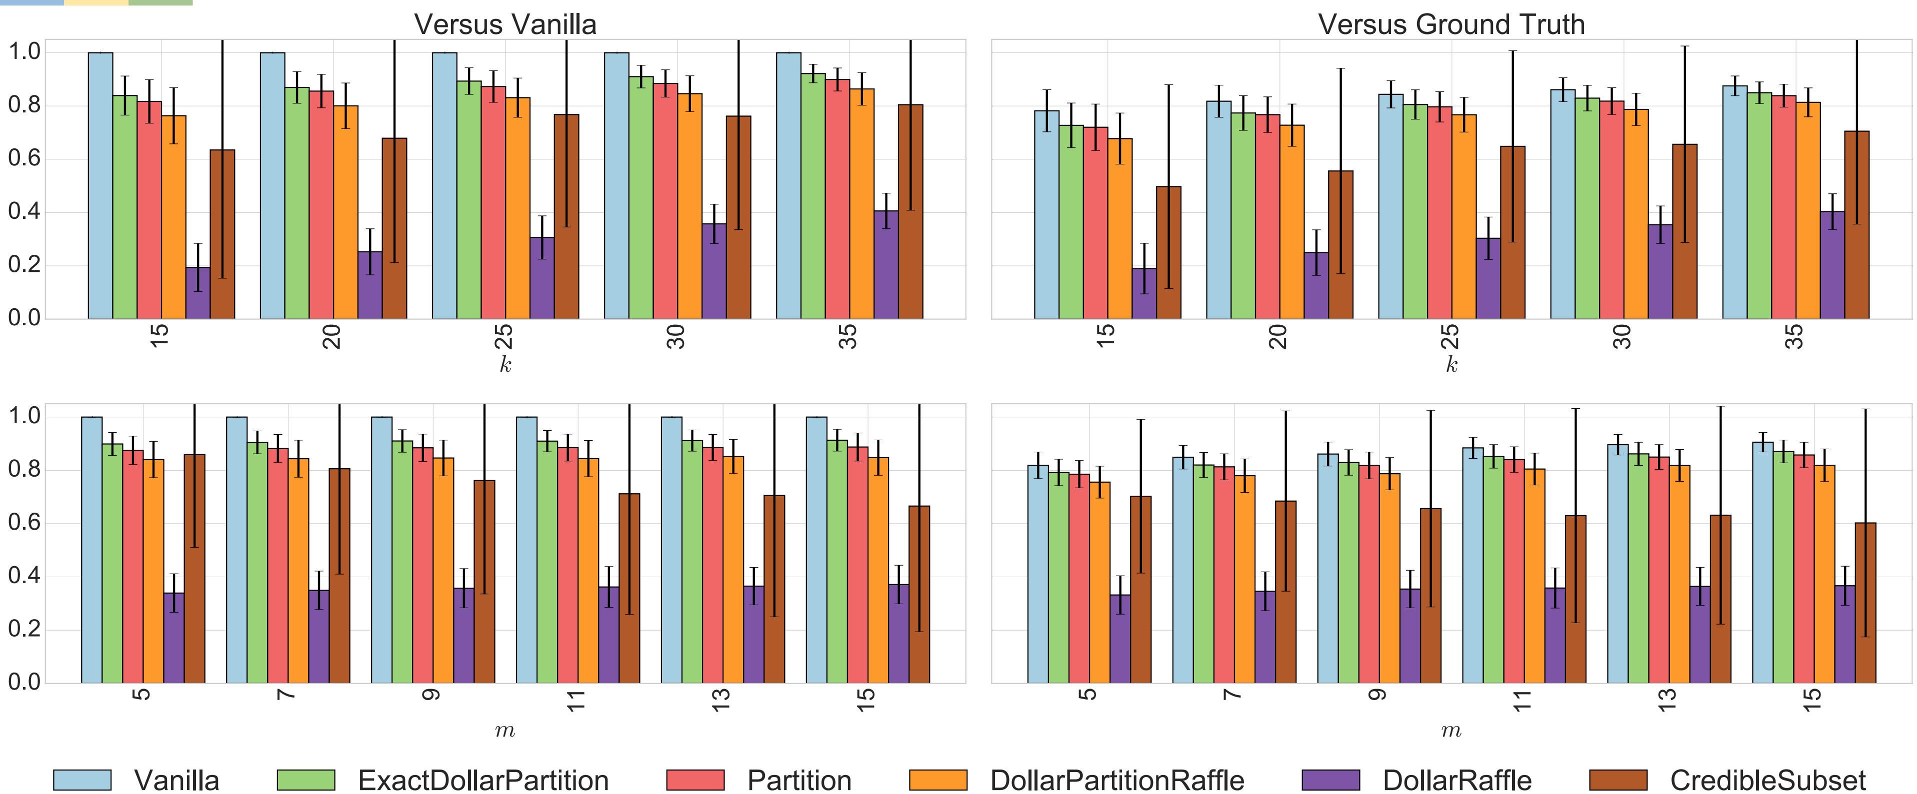

0.5% - 5% better on average, variance 3% - 25% lower.

Results vs. partition

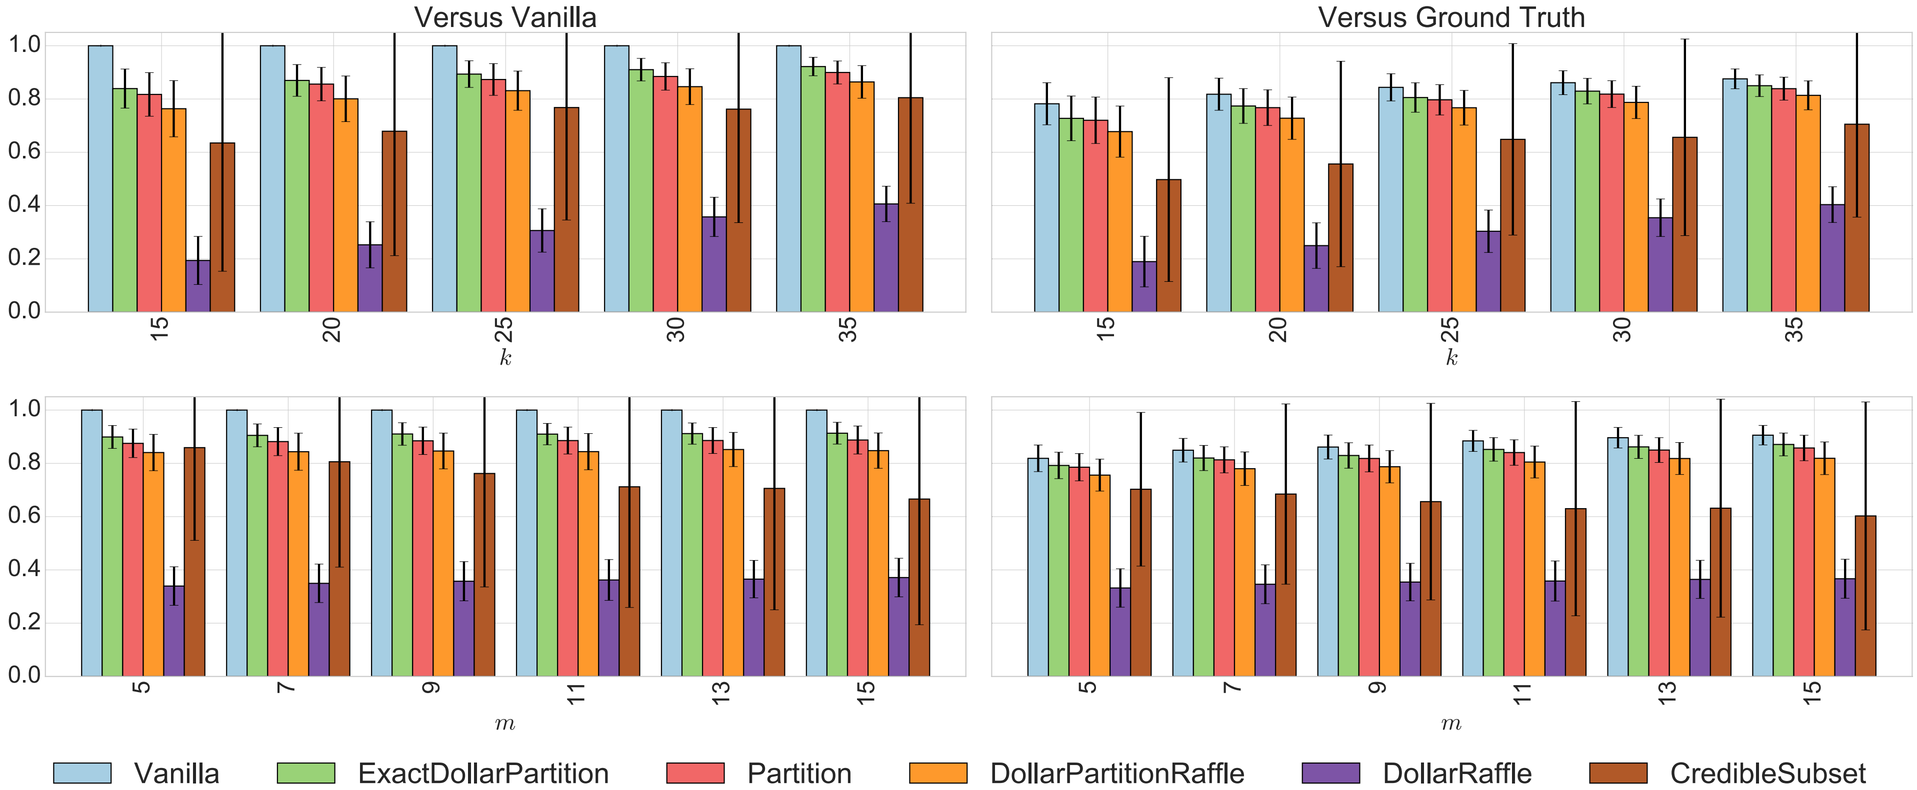

1.5 better proposals on average, 5 better in the worst case.

Results vs. ground truth

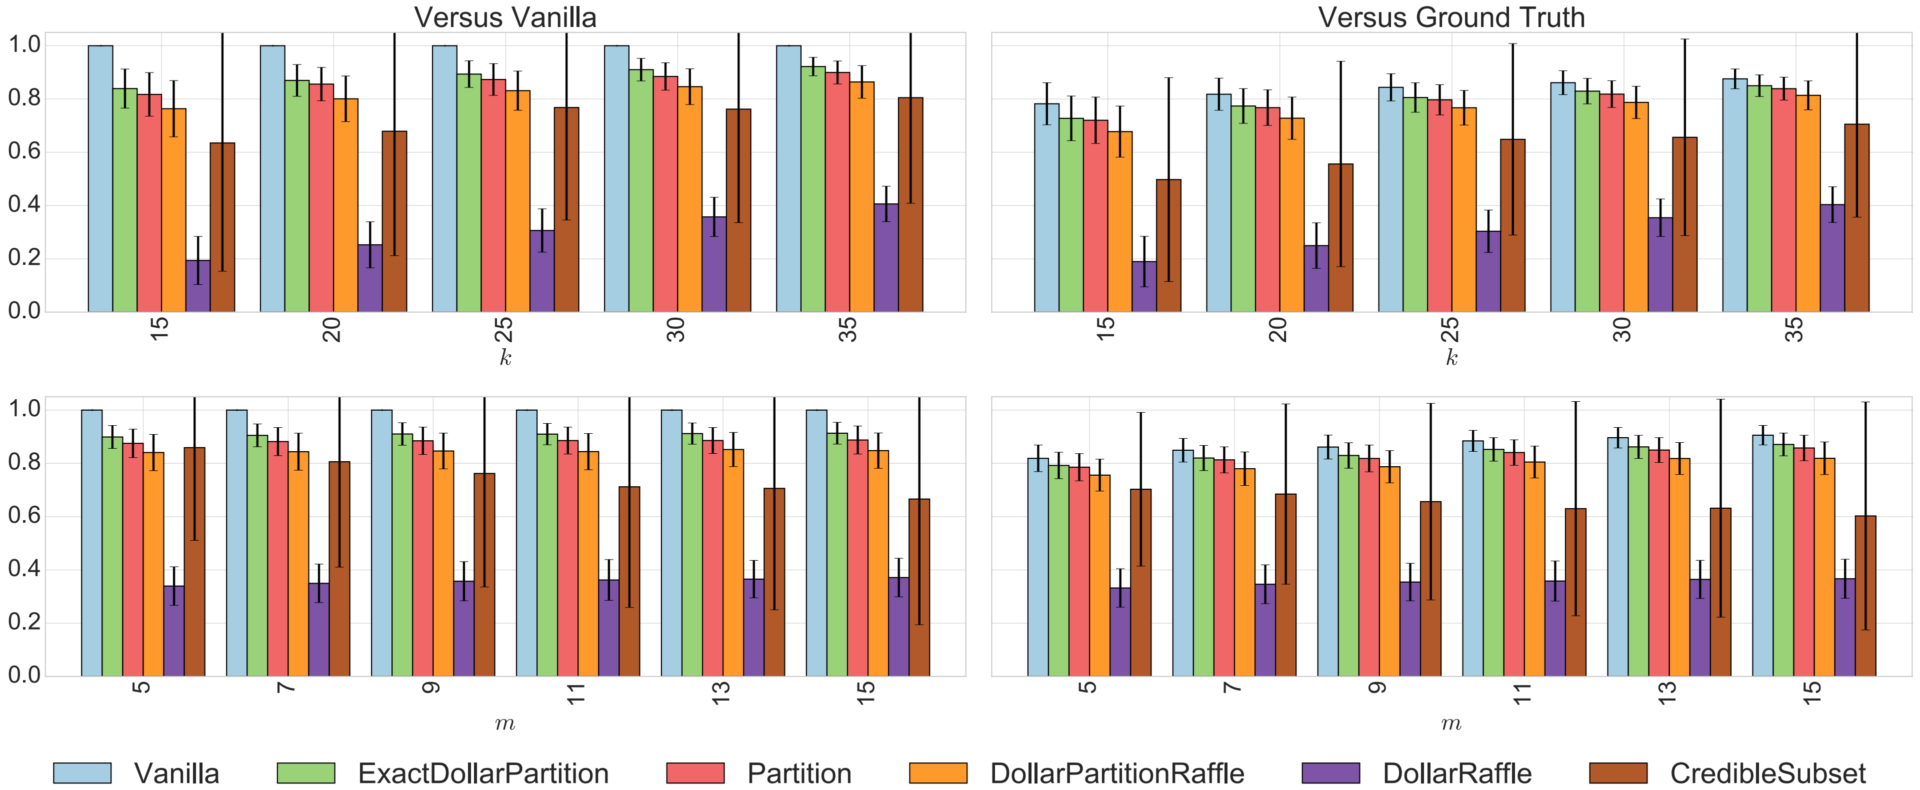

"Cost of strategyproofness" is about 5% of efficiency.

Future work

Implementing in real world cases.

Examining strategyproofness?

More varied comparisons.

How to incentivize work without compromising strategyproofness (too much)?

All simulation code open source and available!