Efficiency Guarantees in Market Design

NICOLE IMMORLICA, MICROSOFT

JOINT WORK WITH B. LUCIER AND G. WEYL

“The greatest risk to man is not
that he aims too high and misses,
but that he aims too low and hits.”

- Michaelangelo

market design.

Allocate scarce resources to agents,

Typical settings:

- FCC: radio spectra to companies
- Business schools: courses to students
- Public schools: schools to students
- Organ transplants: organs to patients
- NRMP: medical residents to residency programs

…to maximize the social good.

market design.

Allocate scarce resources to agents,

… to maximize the social good.

Other considerations:

- Computing time: the market resolves quickly
- Simplicity: it is easy for agents to understand the mechanism and respond to their incentives
- Learning: helps agents find the value of outcomes
- Fairness: agents don’t envy others’ allocations

transferable utility.

Definition: A mechanism is efficient if it maximizes the total utility of the outcome in equilibrium.

Vickrey-Clark-Groves auction (VCG):

1. choose socially optimal allocation according to bids
2. charge prices equal to agents' externality on others

Efficient, yes, but is it computationally quick or simple?

non-transferable utility.

Hard to define “social good” as this requires us to make interpersonal comparisons.

Definition: A mechanism is Pareto efficient if it can’t improve the allocation for some agents without harming others.

Typical mechanisms: deferred acceptance, randomized serial dictatorship, top trading cycles, etc.,

… usually fail ex-ante Pareto efficiency.

approximation.

Paradigm: Guarantee approximate efficiency.

Benefits:

- Sometimes first-best is not achievable
- Relaxing efficiency gives designer more room to optimize other objectives
- Illustrates features of setting that are fundamental for efficiency

transferable utility.

Each agent has a value for each outcome, measured in a common currency.

Goal: maximize total social value.

Definition: A mechanism is efficient if it maximizes the total social value of the allocation in equilibrium.

approximate efficiency.

Definition: A mechanism is α-approximately efficient if the total social value in the worst equilibrium is at least an α fraction of the first-best solution.

Equivalently: A mechanism is α-approximately efficient if the total social value in the worst equilibrium is at least the total social value of the first-best solution in an economy where agents values are scaled by α.

combinatorial allocation.

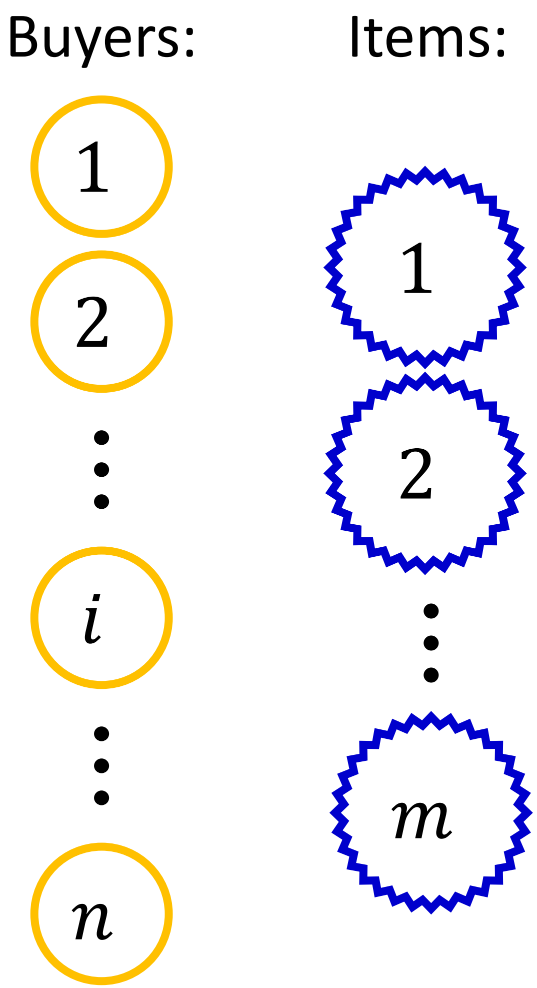

Each buyer has a value for each subset of items,

$v_i(\text{④⑤⑦}) = \$12$

(no assumption on $v_i(\cdot)$)

a simple example.

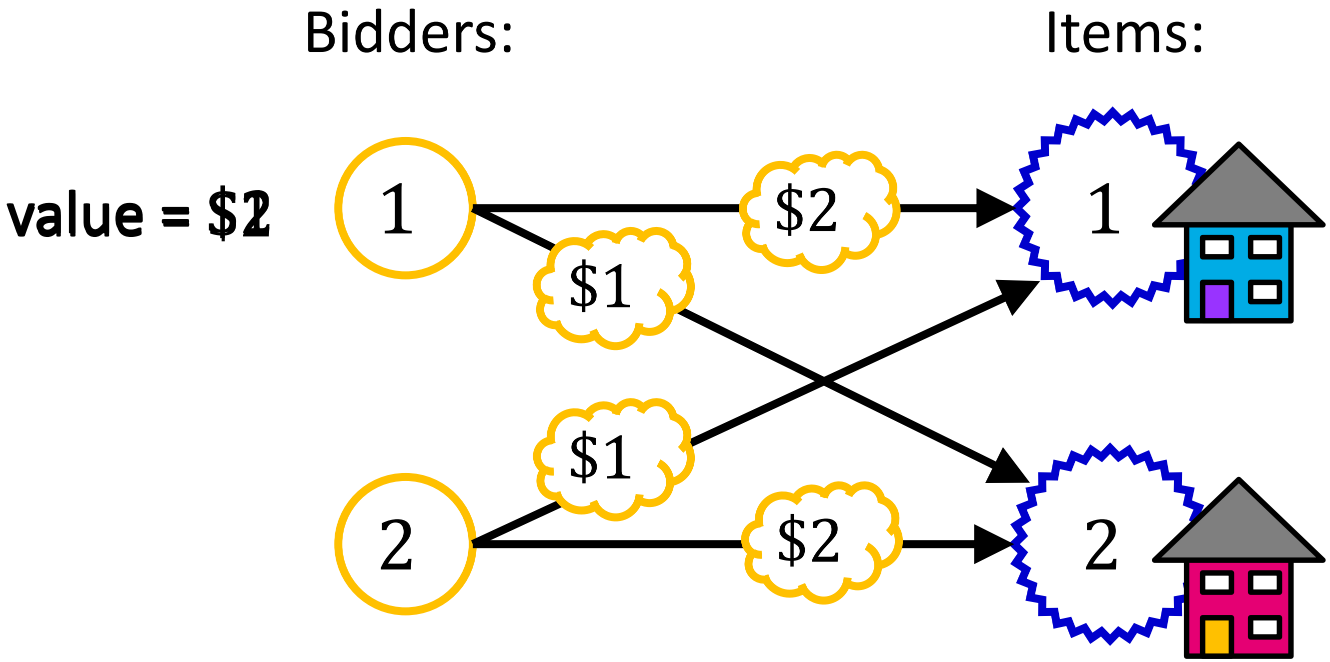

a simple example.

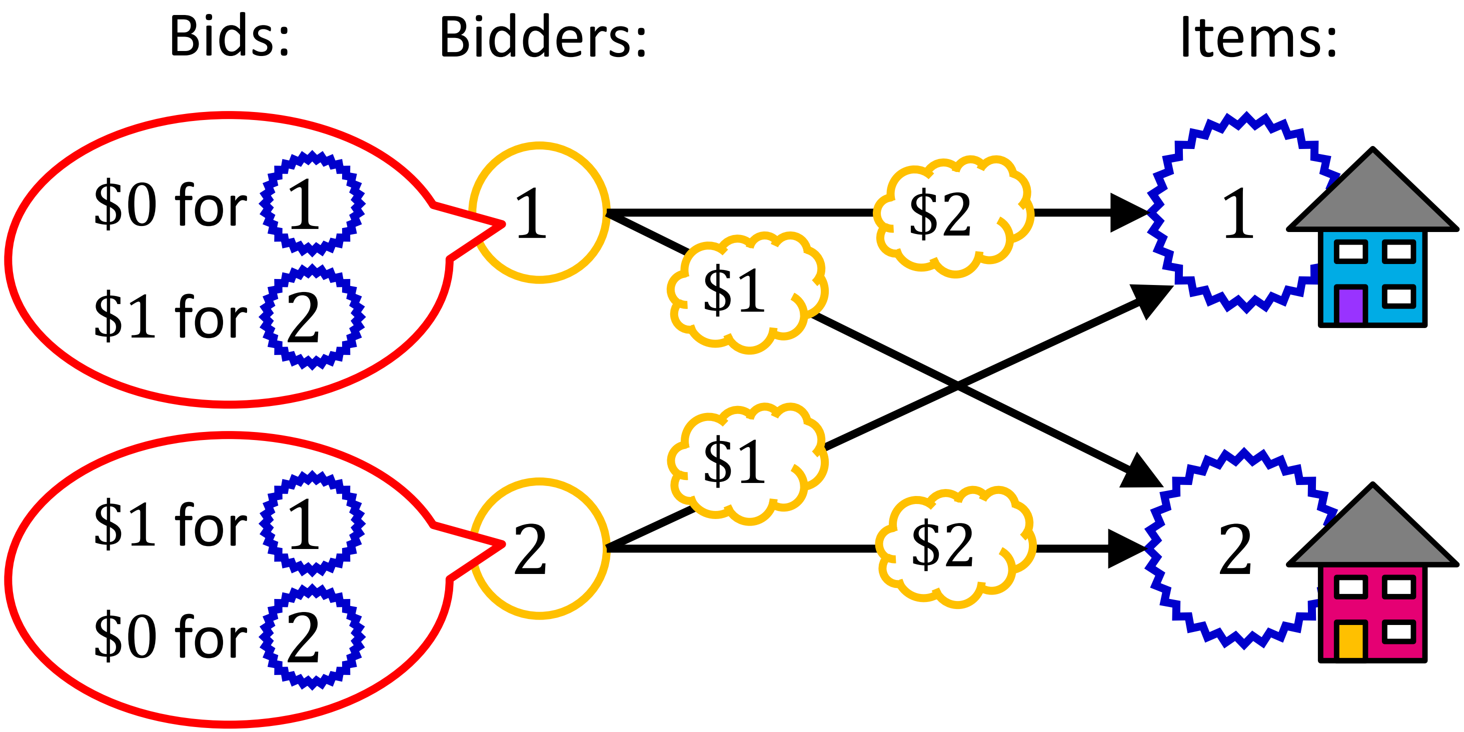

Simultaneous 2nd-price item auctions: bidders place bids on items, each item runs a separate 2nd-price auction with submitted bids.

a simple example.

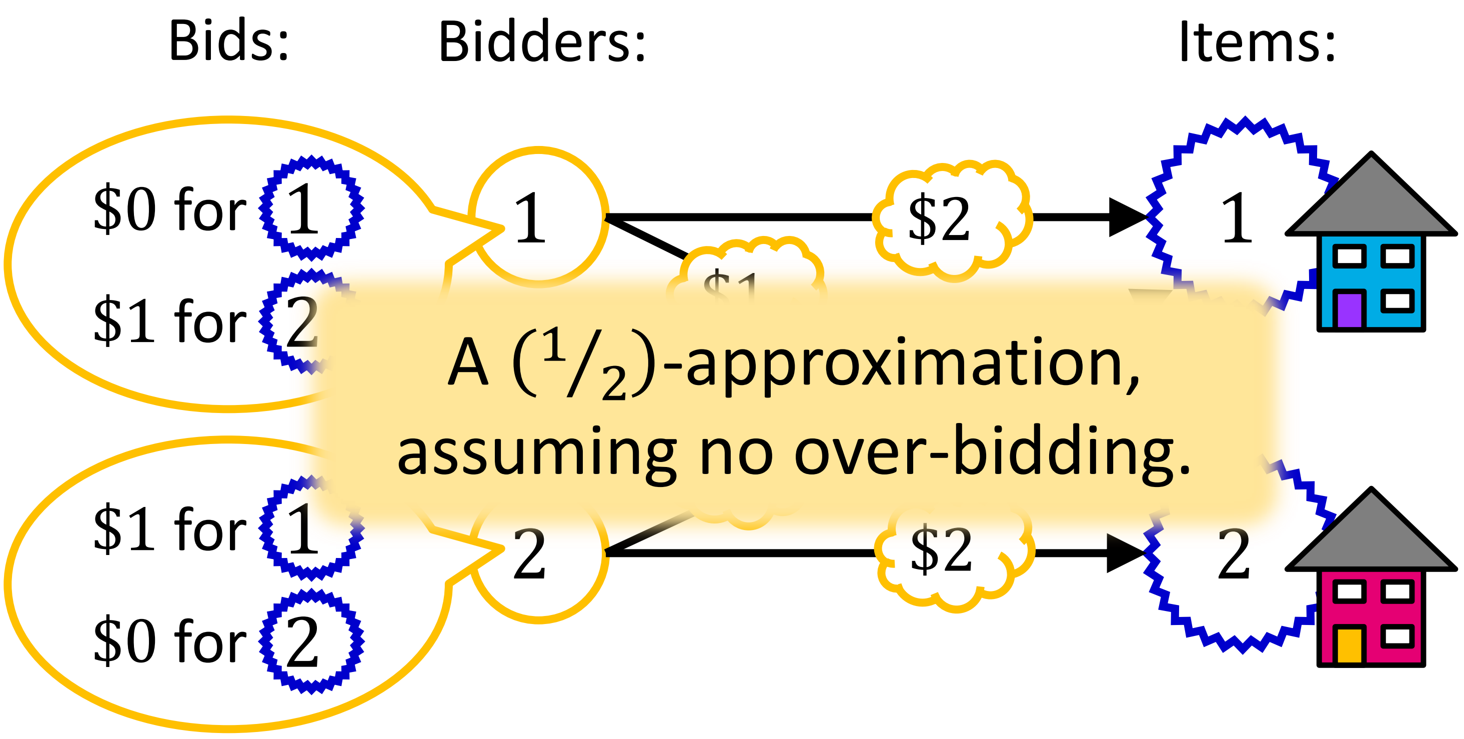

equilibrium outcome: welfare = $2
optimal outcome: welfare = $4

bounding inefficiency.

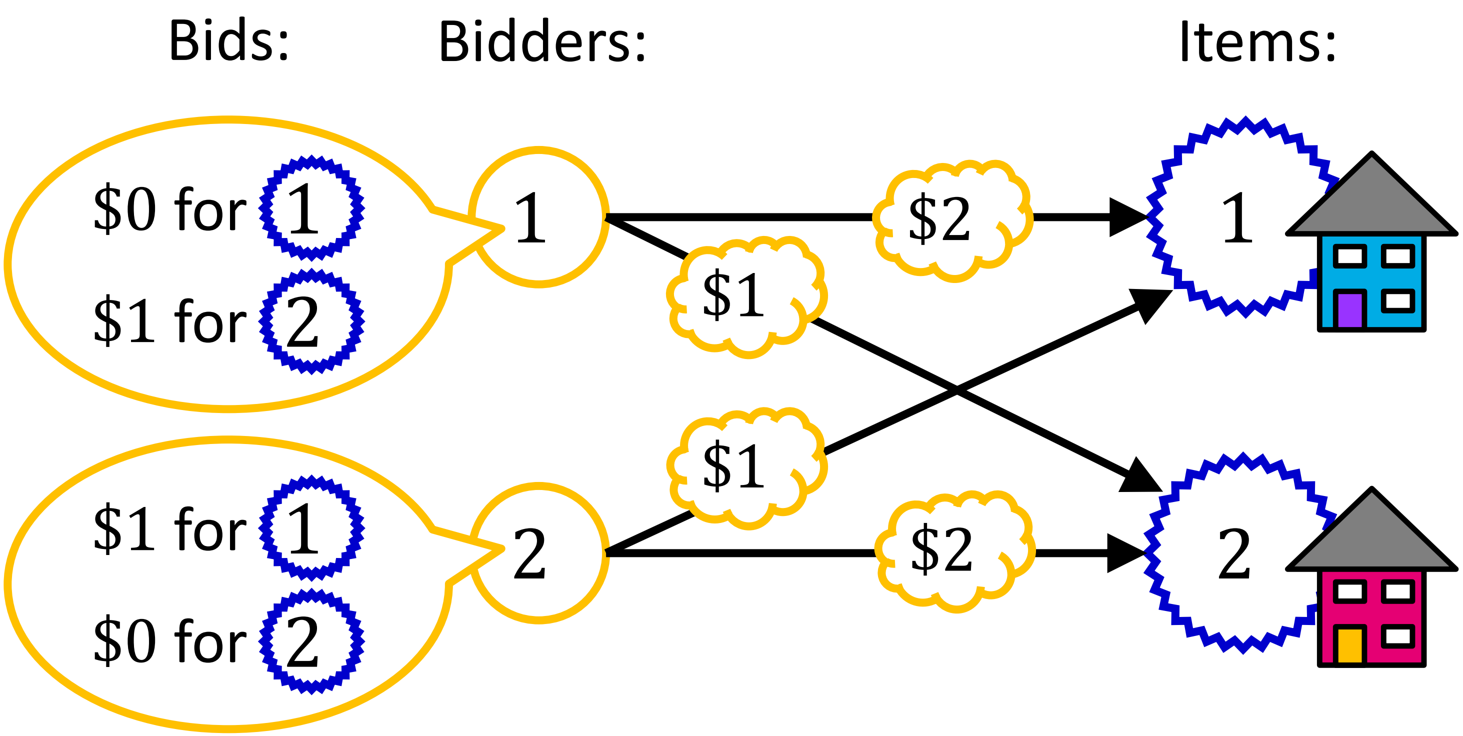

key insight: bidder 1 did not bid on his optimal item because other bidder had a high value for that item.

improvement in large markets.

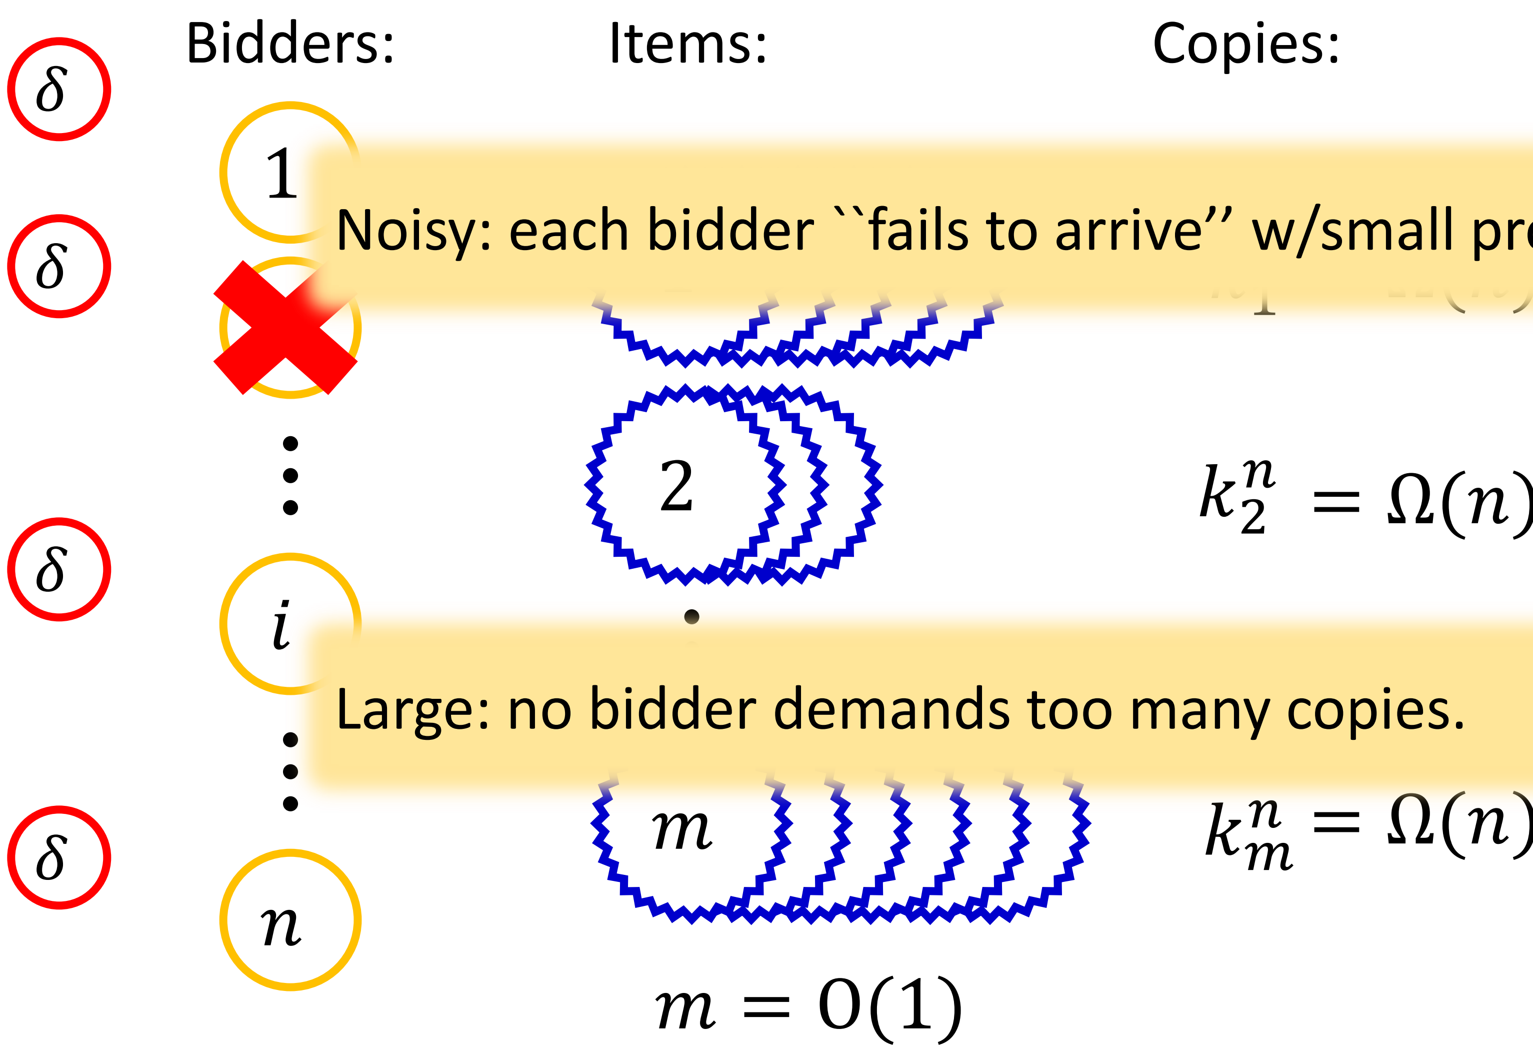

Noisy: each bidder ``fails to arrive’’ w/small prob.

Large: no bidder demands too many copies.

simultaneous uniform price auction.

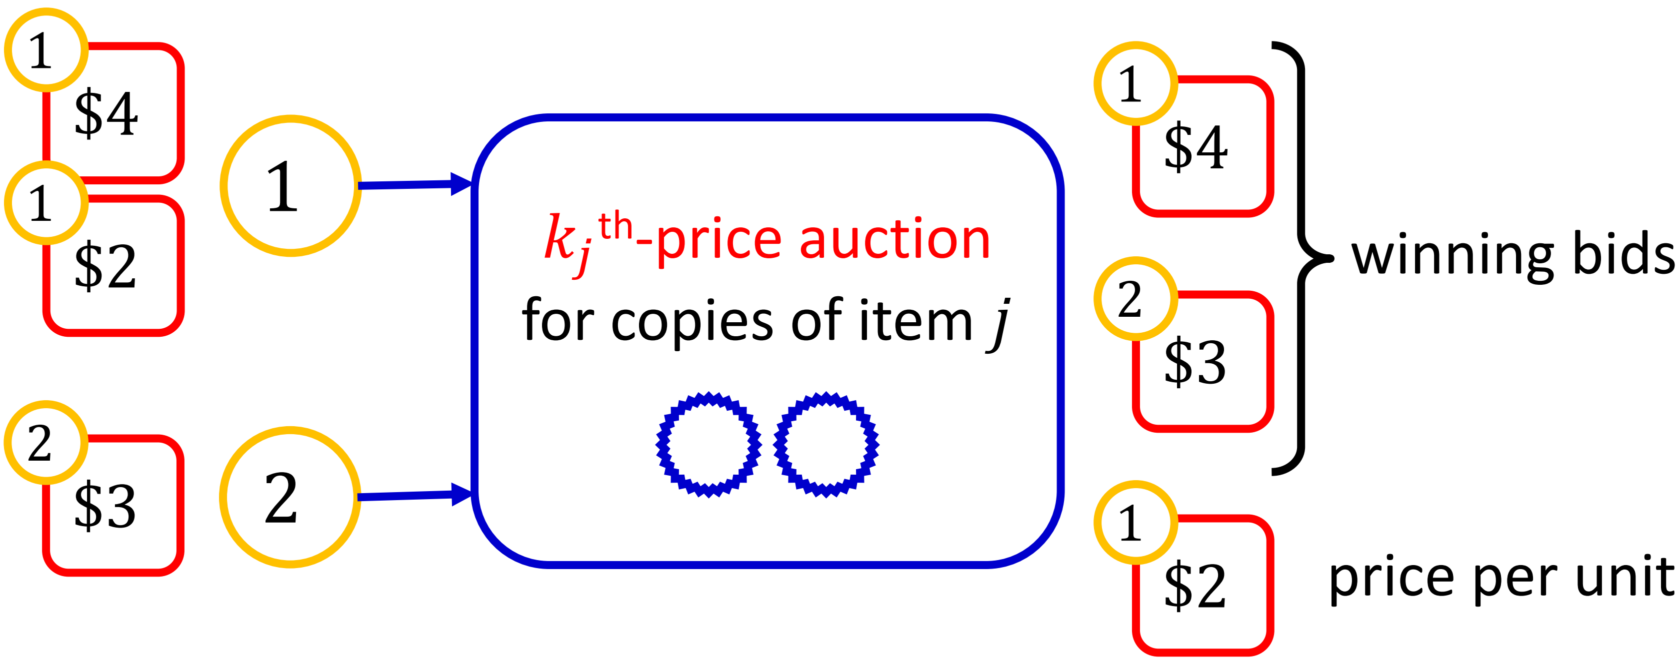

simultaneously for each item j,

1. Each bidder i submits potentially multiple bids.
2. All bids of all bidders for item j are sorted.
3. Top k_j bids win a copy at a price equal to the $(k_j + 1)^{th}$ bid.

efficiency in the large.

Theorem [FILRS'16]. Simultaneous uniform price auctions with uncertain demand approach full welfare in equilibrium at a rate of ${}^{1}/_{\sqrt{n}}$ as market grows large.

Proof Idea. Uses the smoothness [Roughgarden '09], [Syrgkanis-Tardos '13] framework.

If a bidder doesn't receive his optimal allocation,

- it's not profitable for him to bid for this bundle
- hence others' bids are driving up the price
- these bidders must also value this bundle highly

efficiency in the large.

Result holds even though:

- The first-best solution is NP-hard to compute.
- Market-clearing prices might not exist.
- Generically large, noisy markets are not efficient.

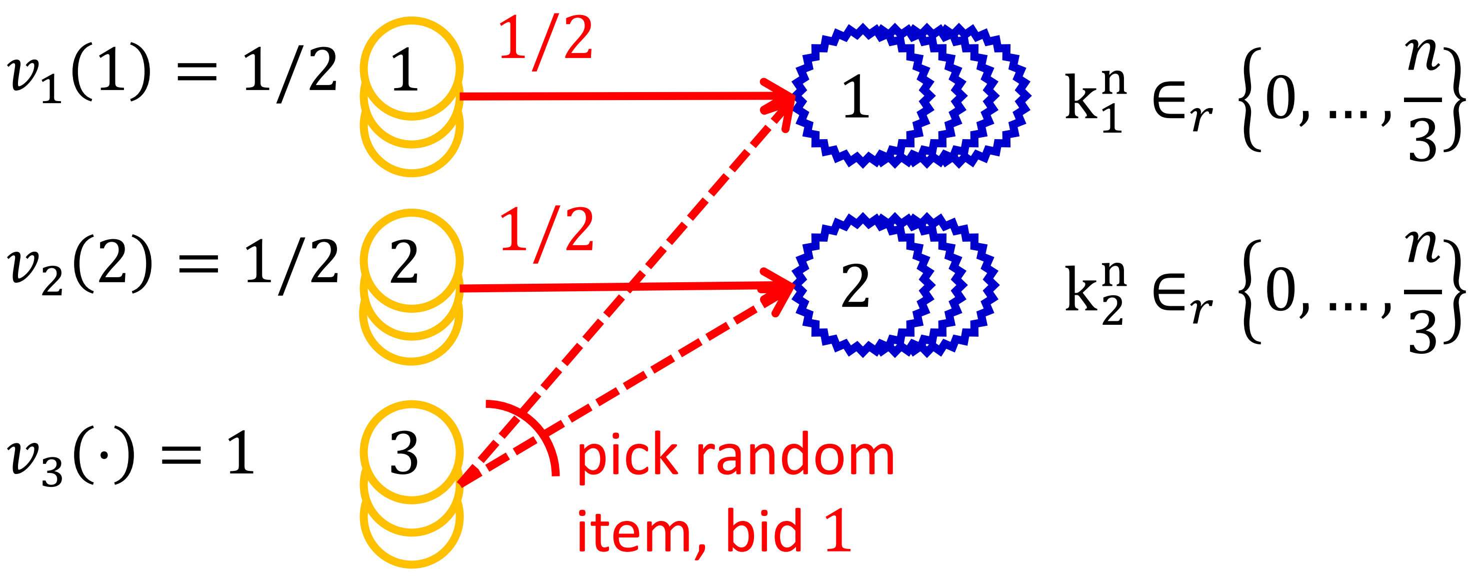

approximate efficiency.

Benefits of approach:

- Guarantees approximate efficiency, bounds degrade nicely as agents approximately optimize
- Illustrates necessity of particular market conditions, demand uncertainty and market size
- Mechanism is simple to describe to participants
- Mechanism is simple to learn how to play

non-transferable utility.

Each agent has a utility for each outcome, measured in incomparable currencies.

Goal: simultaneously maximize all agents' utilities.

Pareto efficiency.

Definition: Given two lotteries, one Pareto dominates the other if each agent's expected utility is weakly higher in that lottery.

Definition: A mechanism is Pareto efficient if the lottery produced by its worst-case equilibrium is not Pareto dominated by any other lottery.

approximate Pareto efficiency.

Definition: A given lottery α-Pareto dominates another if each agent's expected utility is at least an α fraction of his utility in the alternate lottery.

Definition: A mechanism is α-Pareto efficient if it is not α-Pareto dominated by any other lottery.

(Natural analog of transferable-utility approximation.)

an example.

items

agents

	2	1
	1	2

Lottery A:

			utilities
	1	0	2
	0	½	1
	0	½	1

Lottery B half-approximates lottery A.

Lottery B:

			utilities
	0	1	1
	½	0	½
	½	0	½

another example.

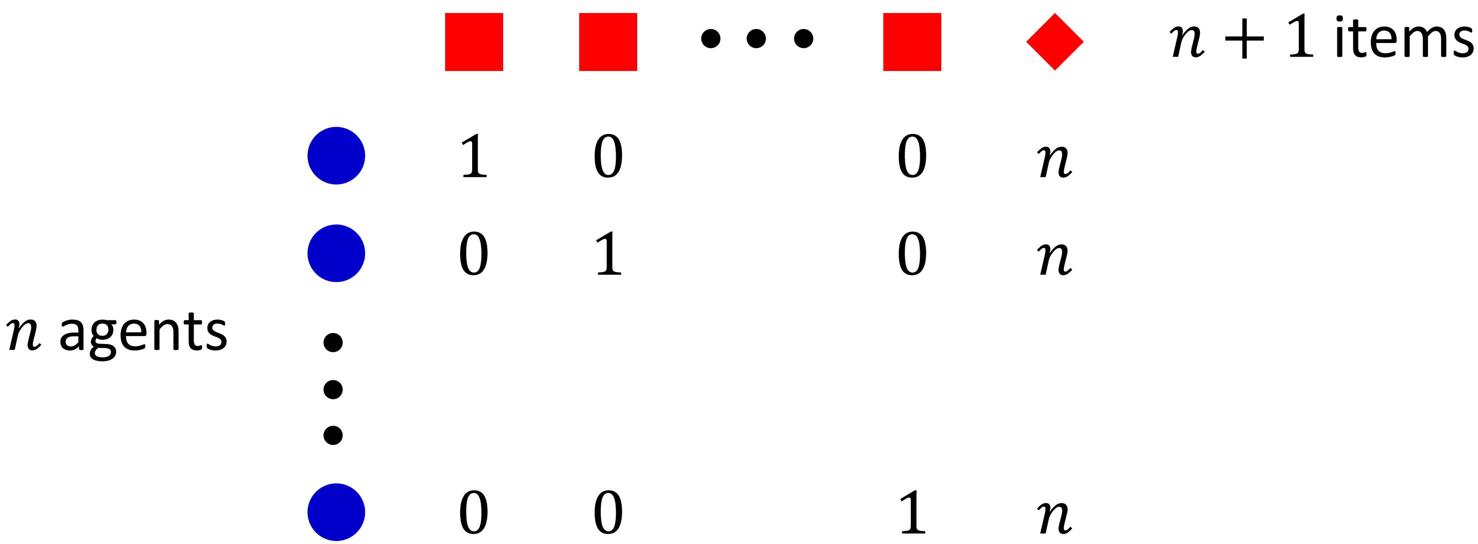

Lottery A: give each agent item ◆ with probability $1/n$.
Lottery B: give agent 1 item ◆, agent $i > 1$ item i.

approximate Pareto efficiency.

Alternate definition: No other lottery improves some agent by an α-factor while not harming others.

Poor choice:

- Requires designer to designate special agent.
- Does not distinguish between mechanisms.

towards efficiency.

Key Element: In non-transferable utility settings, probability of allocation acts as a price, allowing agents to express strength of preference, yielding more efficient outcomes.

"The question, love, is whether you want me enough to take the risk."

- some romance novel

(disclaimer: I don't read romance novels, … no really)

school choice.

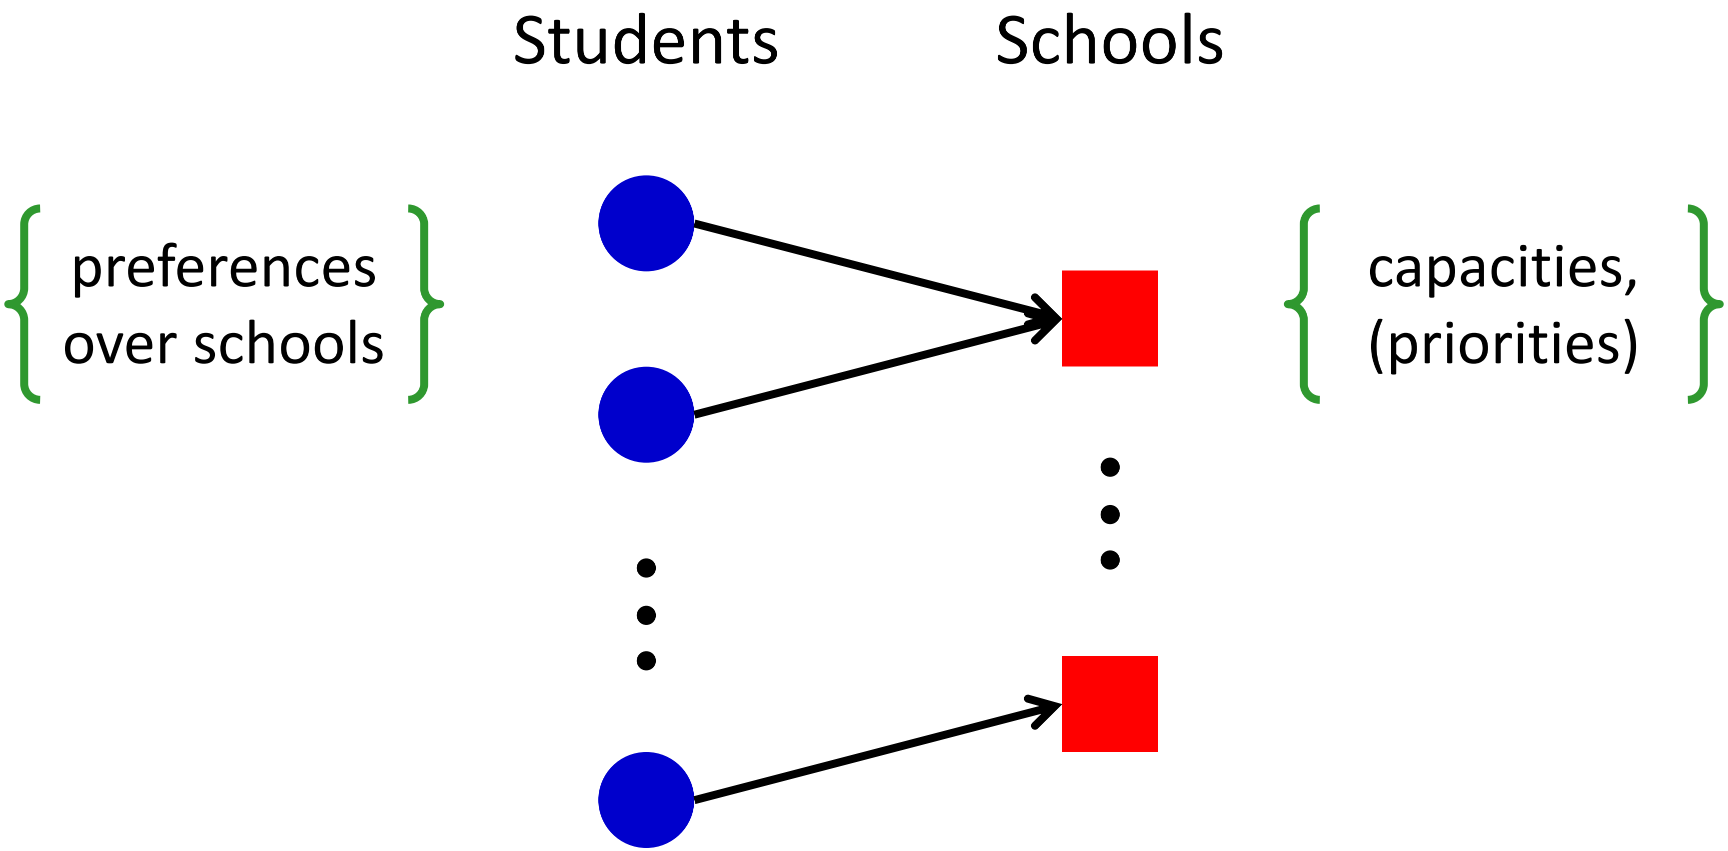

school choice.

Student i has value v_{ij} for school j.

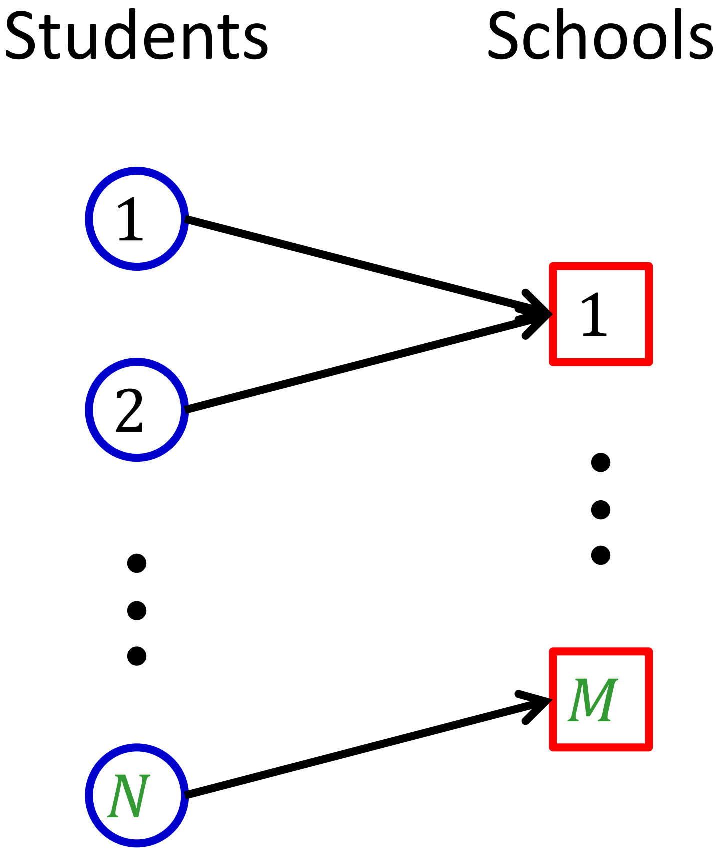

School j has capacity c_j, no priority.

school choice.

Outcome: an assignment of students to schools that respects capacities.

Desirable properties

- Pareto efficiency: no way to improve allocation for some student without harming other students
- Fairness (stability): no “justified envy” wrt priorities
- Simplicity: no incentive to lie (truthfulness), “can be explained to a fifth-grader” (Peng Shi)

impossibility.

Theorem [Kesten, 2010]: There is no Pareto efficient and truthful mechanism that selects the Pareto efficient and stable match whenever one exists.

"Two out of three ain't bad."

- Meat Loaf

deferred acceptance.

Students

Schools

Students iteratively apply to schools in order of preference.

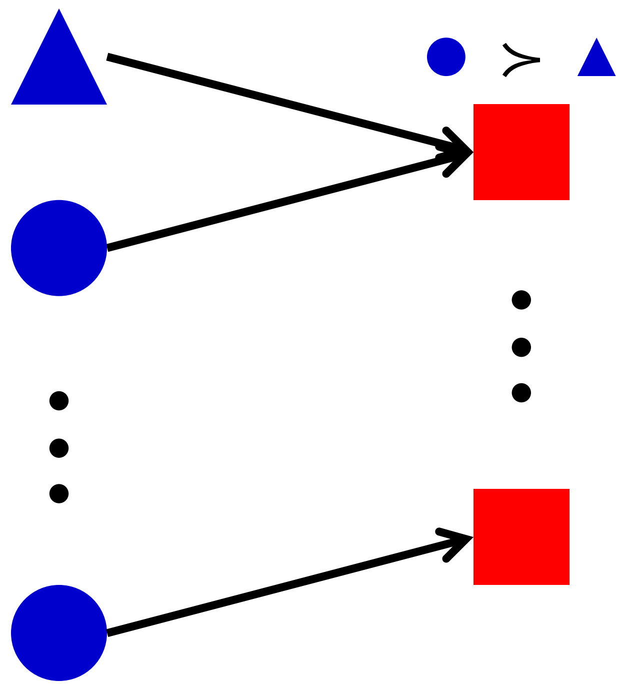

Schools tentatively accept students in order of priority.

(random tie-breaking)

deferred acceptance.

Deferred acceptance is truthful and stable,

… but not Pareto efficient.

		s_1	s_2	s_3	(no priorities)
(two students)	▲	$10 \times \frac{1}{3}$	$2 \times \frac{1}{3}$	$0 \times \frac{1}{3}$	$= 4$
(one student)	●	$7 \times \frac{1}{3}$	$5 \times \frac{1}{3}$	$0 \times \frac{1}{3}$	$= 4$

Pareto-improvement: trade ●'s s_1 shares for s_2.

deferred acceptance.

Example: Approximation factor can be arbitrarily bad.

schools: n schools, 2^i seats at school i

students:
n types,
2^n per type

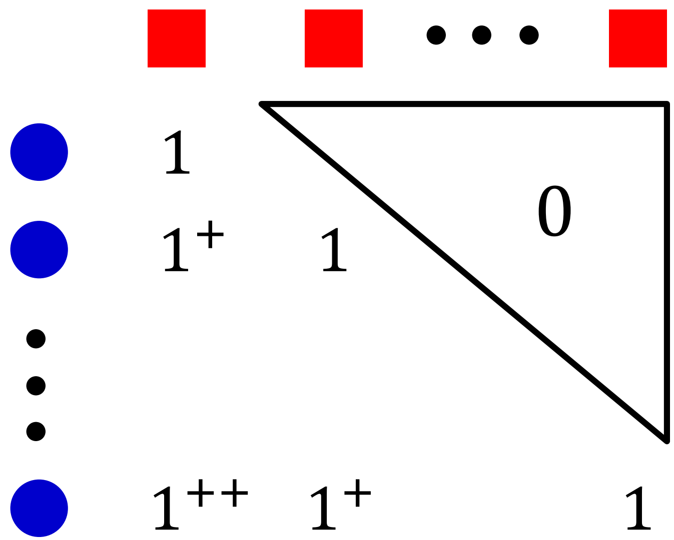

utility of type i:

$$\frac{1}{n2^n} \cdot \sum_{j=1}^{i} 2^j \leq \frac{2^{i+1}}{n2^n}$$

alternative: assign items of type i randomly to agents of type i,

utility of type i is $\frac{2^i}{2^n}$, a factor $\frac{n}{2}$ better than deferred acceptance.

Boston mechanism.

Students Schools

Students iteratively apply to schools in order of preference.

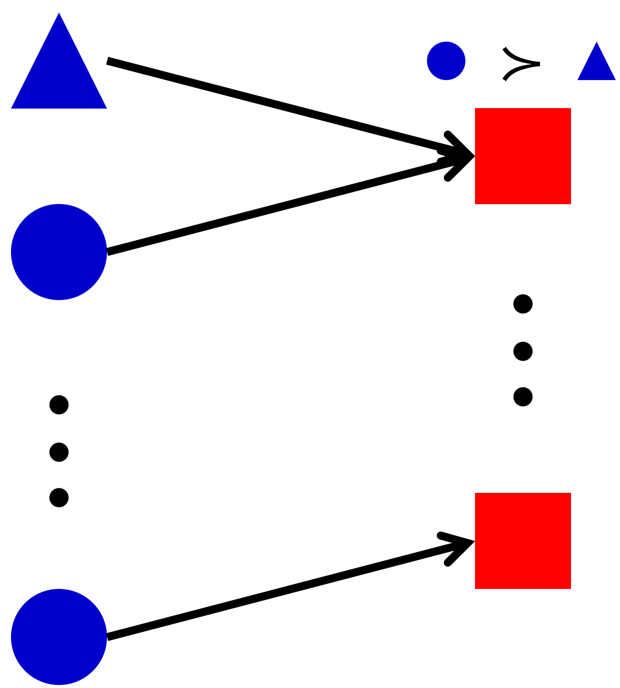

Schools irrevocably accept students in order of priority.

(random tie-breaking)

Boston mechanism.

Boston mechanism is not truthful or stable,

… but let's agents express strength of preference.

		s_1	s_2	s_3	(no priorities)
(two students)	▲	$10 \times \frac{1}{2}$	2×0	$0 \times \frac{1}{2}$	$= 5$
(one student)	●	7×0	5×1	0×0	$= 5$

efficiency guarantees.

Theorem: When students' preferences are aligned and schools have no priorities, the Boston mechanism Pareto dominates deferred acceptance.

Theorem: In large markets, when preferences are sufficiently aligned, a mashup of the Boston mechanism and deferred acceptance called the choice-augmented deferred acceptance mechanism is Pareto efficient.

[Abdulkadiroglu, Che, and Yasuda; 2011, 2015]

efficiency guarantees.

Example: Boston mechanism not fully efficient.

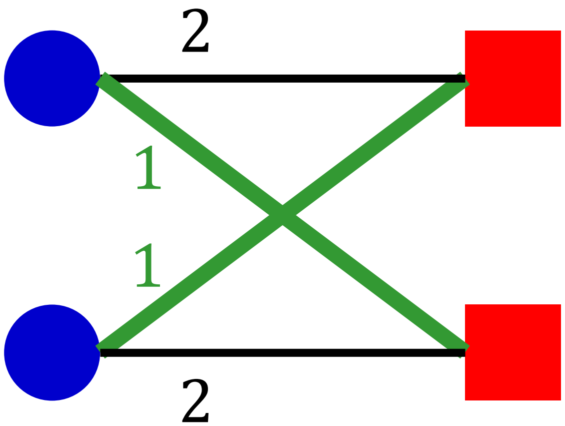

Agents fail to coordinate on good equilibrium.
Agent’s choice changes probability of item allocation.

single-ticket raffle.

Mechanism:

Each item has a bucket.

Each agent has one ticket to put in bucket of his choice.

1. In arbitrary order, draw c_j tickets from bucket j.
2. Assign agent to item if his ticket drawn from its bucket, otherwise agent gets no item.

(Like just running 1st iteration of Boston mechanism.)

approximate efficiency.

Lemma: Raffle mechanism is $({}^{1}/_{2})$-Pareto efficient, and fully efficient in continuum, if all items are congested.

Proof: Fix equilibrium π, let n_j be # tickets in bucket j.

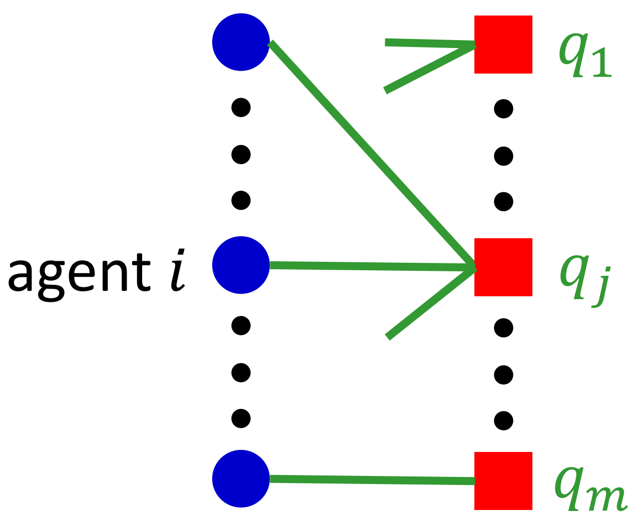

Congestion ratio: $\mathrm{q_j} = \min\left(\frac{\mathrm{c_j}}{\mathrm{n_j}}, 1\right)$, interpret as a "unit" of item j.

Equilibrium condition for i chosing j:

$$v_{ij} q_j \geq v_{ik} \cdot \frac{c_k}{n_k + 1} \geq \frac{v_{ik} q_k}{2}$$

approximate efficiency.

Theorem: Raffle mechanism is $(^1/_2)$-Pareto efficient, and fully efficient in continuum, if all items are congested.

Proof, continued: Given equilibrium π,

- Let ρ be an alternate lottery's assignment of units.
- Let $x_i = \sum_j \rho_{ij} q_j$ be total # units of agent under ρ.
- Note if $x_i < 1$ for some i, then π gives i (at least half) his value in ρ by equilibrium condition.
- But since items congested, infeasible to have $x_i \geq 1$ for all agents with a strict inequality for some agent.

Boston mechanism.

Reinterpret as a raffle:

Each agent has a primary ticket, a secondary ticket, etc.

Each item has primary bucket, secondary bucket, etc.

1. Agents deposit primary tickets in primary buckets, secondary tickets in secondary buckets, etc.
2. Draw tickets from buckets, starting with primary buckets, then secondary, etc.
3. Assign agent to item if his ticket drawn from one of its buckets, otherwise agent gets no item.

extension to Boston mechanism.

Theorem: Boston mechanism is $({}^{1}/_{2})$-Pareto efficient in continuum.

Proof Idea: Consider equilibrium π, alternate lottery ρ.

1. Uncongested items, i.e., those with $q_j \geq 1$ for primary buckets, can contribute at most half the value of ρ.
2. Can translate agent values by outside option from other tickets to induces a single-ticket raffle game and so π is fully efficient on congested items.

extension to Boston mechanism.

Theorem: Boston mechanism is $(^1/_2)$-Pareto efficient in continuum.

Work-in-progress: Use properties of secondary, etc., ticket distributions to tighten approximation.

raffle mechanisms.

Outcome is:

- Approximately Pareto efficient
- Fair, no agent envies another's lottery

Is this simple?

- Easy to describe
- Agents must reason strategically, but can use historical congestions as prices
- Aids information acquisition if it is costly for an agent to learn her values

conclusion.

Approximation is a powerful tool for designing good mechanisms when first-best is infeasible.

Approach gives designer room to optimize secondary goals while maintaining a worst-case guarantee.

Can highlight importance of market assumptions.