

Random Assignment Mechanisms

November 8, 2011

Allocating Objects

- ▶ Assign objects to agents
 - ▶ Student placement into schools - where schools have no preferences
 - ▶ House allocation in colleges
- ▶ Constraints:
 - ▶ No Monetary transfer
 - ▶ Only ordinal preferences can be used: Ask agents to rank various goods.
- ▶ Two major mechanisms:
 - ▶ Random Priority mechanism (RP)
 - ▶ Probabilistic Serial Mechanism (PS)

The Model

- ▶ There is a fixed set of good O that can be distributed among agents, plus the null good - receiving no good, ϕ
- ▶ Each agent i has strict preferences over goods including the null good: $O \cup \{\phi\}$
- ▶ A randomn assignment is a matrix $P = [P_{ia}]_{i,a}$ where P_{ia} is the probability that agent i obtains good a .

The Random Priority Mechanism

Random Priority (RP)

- ▶ Draw each possible ordering of the agents with equal probability
- ▶ The first agent receives her most preferred unit, the next agent his most preferred unit among the remaining units, and so on

Random Priority is:

- ▶ Easy to implement
- ▶ Strategy-Proof
- ▶ Widely used in practice

RP is Inefficient

Suppose there are two good $O = \{a, b\}$ with one copy each and agents $N = \{1, 2, 3, 4\}$.

- ▶ 1 and 2's preferences: a, b, ϕ
- ▶ 3 and 4's preferences: b, a, ϕ

The random assignment under RP:

	Good a	Good b	Good ϕ
Agents 1 and 2	5/12	1/12	1/2
Agents 3 and 4	1/12	5/12	1/2

What is the problem?

The random assignment under RP:

	Good a	Good b	Good ϕ
Agents 1 and 2	$5/12$	$1/12$	$1/2$
Agents 3 and 4	$1/12$	$5/12$	$1/2$

Everyone prefers

	Good a	Good b	Good ϕ
Agents 1 and 2	$1/2$	0	$1/2$
Agents 3 and 4	0	$1/2$	$1/2$

Ordinal Efficiency

- ▶ A random assignment P (first-order) stochastically dominates another random assignment P' if for every agent

$$\Pr[i \text{ gets } a \text{ or a more preferred item under } P]$$

$$\geq \Pr[i \text{ gets } a \text{ or a more preferred item under } P'], \text{ for all } i, a$$

with strict inequality for at least one pair i, a .

- ▶ A random assignment is ordinally efficient if it is not stochastically dominated by any other random assignment.
- ▶ In environments where only ordinal preferences can be used, ordinal efficiency is a natural efficiency concept.
- ▶ RP may result in ordinally inefficient random assignments (see the last example).

Probabilistic Serial Mechanism

Bogomolnaia and Moulin (2001) define PS based on (cake) eating algorithms

1. Imagine each good is divisible, and is divided into probability shares of that good
2. Suppose there is a time interval $[0, 1]$
3. Each agent “eats” the good she prefers the most among all goods that have not been eaten yet with speed one
4. At time $t = 1$, each agent is endowed with probability shares
5. The PS assignment is the resulting profile of shares.

PS mechanism example

Consider the example from before: two good $O = \{a, b\}$ with one copy each and agents $N = \{1, 2, 3, 4\}$.

- ▶ 1 and 2's preferences: a, b, ϕ
- ▶ 3 and 4's preferences: b, a, ϕ

Compute the PS assignment

- ▶ $t=0$: Agents 1 and 2 start eating a , agents 3 and 4 start eating b
- ▶ $t=1/2$: Goods a and b are eaten away. No goods remain, only the null good ϕ .

The resulting assignment:

	Good a	Good b	Good ϕ
Agents 1 and 2	5/12	1/12	1/2
Agents 3 and 4	1/12	5/12	1/2

This assignment is ordinally efficient.

PS is ordinally efficient

Theorem:

For any reported preferences, the PS mechanism produces an ordinally efficient assignment with respect to the reported preferences.

Intuition:

At each moment, everyone is eating her favorite available good
So, if an agent would prefer a probability share of a different good rather than the one she is eating from right now, then that preferred good must have already been eaten away.

PS is fair

To think about fairness of an outcome, we can think back to stability in two-sided mechanisms.

A random assignment is envy free if everyone prefers his or her assignment to the assignment of anyone else (i.e. his or her assignment stochastically dominates the assignments of others).

Theorem:

For any reported preferences, the PS mechanism produces an envy-free assignment with respect to the reported preferences.

Intuition: Once more: at any point in time anyone eats from their most preferred good. So, anyone eats a good they prefer to the good others are eating, so in the end no one envies the random assignment of anyone else.

PS is not strategy-proof

Consider two good $O = \{a, b\}$ with one copy each and agents $N = \{1, 2, 3, 4\}$.

- ▶ 1's preferences: a, b, ϕ
- ▶ 2's preferences : a, ϕ
- ▶ 3 and 4's preferences: b, ϕ

If 1 reports her preferences truthfully:

	Good a	Good b	Good ϕ
Agents 1 and 2	1/2	0	1/2
Agents 3 and 4	0	1/2	1/2

If 1 reports b, a, ϕ .

	Good a	Good b	Good ϕ
Agent 1	1/3	1/3	1/3
Agent 2	2/3	0	1/3
Agents 3 and 4	0	1/3	2/3

For some cardinal preferences, 1 is made better off.

Comparing PS and RP

- ▶ PS has better efficiency properties (ordinal efficient)
- ▶ RP has better incentive properties: dominant strategy

Theorem: Bogomolnaia and Moulin 2001

There is no mechanism that is ordinally efficient, strategy-proof and symmetric (i.e., two agents who report the same preferences get the same random assignment.)

How do the outcomes fare in data?

Pathak 2008 uses NYC data to compare RP and PS.

NYC's supplementary round: currently use RP. Note that RP is strategy-proof, so we can expect that the submitted preferences are truthful.

Pathak notes that the difference in outcomes is small: about 4,999 students (out of 8,255) receive their top choice from RP, while 5,016 students receive their top choice from PS, a difference of 0.3%.

Based on that, he supported RP over PS.

Further Questions for PS

What happens in large markets?

What is the right efficiency measure?

PS Ordinal efficiency compared to Boston (or rank efficiency) in a simple example.

- ▶ Recall the example where we have 2 schools, an Art school A , and a science school, S .
- ▶ Each school has one seat. There are 3 students, i_1, \dots, i_3 . Each student is equally likely to be either an art or science student, where the art student prefers the art over the science school and the science student prefer the science over the art school.
- ▶ What are the expected outcomes of Boston and PS

Boston	1st choice	2nd choice	No School
Student 1	1	0	0
Student 2	1/2	1/4	1/4
Student 3	1/4	0	3/4
Expected	$1/2 + 1/12$	$1/12$	$1/3$

In Boston, whenever there is an art student, the art school is filled with an art student, and likewise for the Science school

What happens under PS?

Suppose there are 2 art and one science student.

- ▶ $t=0..1/2$: the 2 art students eat the art school, the science student eats $1/2$ of the science school
- ▶ $t=1/2..$: all 3 students eat the science school.
- ▶ Science student has a $1/2 + 1/6 = 2/3$ chance to enter the science school
- ▶ Under DA Science students had a $2/3$ chance to enter the science school.

Large Markets Results

Theorem (Kojima and Manea 2007)

Fix agent i 's utility function u_i , and assume u_i represents strict preferences. There is a finite bound M such that, if $q_a \geq M$ for all $a \in O$, then truthtelling is a dominant strategy for i under PS. The conclusion holds no matter how many other agents are participating in the market.

- ▶ Remark Truthtelling is an exact dominant strategy in finitely large markets.
- ▶ How small can M be? Consider a school context, where a student finds only 10 schools acceptable, and her utility difference between any two consecutively ranked schools is constant. Then truthtelling is a dominant strategy for her in PS if each school has at least 18 seats.

Intuition for the proof:

Manipulations have two effects:

1. given the same set of available objects, reporting false preferences may prevent the agent from eating his most preferred available object;
2. reporting false preferences can affect expiration dates of each good.

(1) always hurts the manipulating agent, while (2) can benefit the agent.

Intuitively, the effect (2) becomes small as the market becomes large.

A nontrivial part of the formal proof is that (2) becomes very small relative to (1) when the copies of each object type becomes large, so the agents hurt themselves in total.

Difference between PS and RP in large markets

What is the difference between PS and RP as the market becomes large.

How should the market become large? Keep the number of items fixed and have more and more copies, or increase the number of items of a given size?

Che and Kojima 2009:

a q -economy consists of

- ▶ q copies of each (real) good and infinite copies of ϕ , and a set of agents:
- ▶ Assume that

$$\frac{\text{(number of agents with preference } \pi \text{ in } q\text{-economy)}}{q}$$

converges as $q \rightarrow \infty$ for every preference π (the limit can be zero).

Theorem (Che and Kojima (2009))

Fix the set of types of goods. The random assignments in RP and PS converge to each other as $q \rightarrow \infty$. Formally,

$$\lim_{q \rightarrow \infty} \max_{\pi, a} RP_a^q(\pi) - PS_a^q(\pi) = 0, \text{ where}$$

$RP_a^q(\pi) := \Pr[\text{agents with pref } \pi \text{ get } a \text{ in a } q\text{-economy in } RP]$

$PS_a^q(\pi) := \Pr[\text{agents with pref } \pi \text{ get } a \text{ in a } q\text{-economy in } PS]$

Intuition

- ▶ In PS, the random assignment is completely pinned down by the expiration dates of the goods. The expiration date T_a^q of good a is the time at which good a is completely eaten away.
- ▶ The probability that an agent receives good a is equal to the duration at which he is consuming good a , so it is

$$\max\{T_a^q - \max\{T_b^q \mid b \text{ is preferred to } a\}, 0\}$$

- ▶ Proof idea: Find RP-analogues of expiration dates, and show that they converge to expiration dates in PS, in probability.

Alternative formulation of RP.

1. Each agent draws a number iid uniformly distributed in $[0; 1]$.
 2. The agent with the smallest draw receives her favorite good, and so on.
- ▶ Given the realized draws, the cutoff \hat{T}_a^q of good a under RP is the draw of the agent who receives the last copy of a .
 - ▶ Since random draws are uniform over $[0; 1]$, an agent will receive good a with probability

$$E[\max\{\hat{T}_a^q - \max\{\hat{T}_b^q | b \text{ is preferred to } a\}, 0\}]$$

- ▶ Fuhito and Che show that cutoffs of RP converge to expiration dates of PS (in probability).
- ▶ They are different in general: In PS, a good is consumed proportionately to the number of agents who like it: In RP, a good may be consumed disproportionately to the number of agents who like it because of the randomness of draws.
- ▶ For RP in large markets, the law of large numbers kicks in: with a very high probability, a good is consumed almost proportionately to the number of agents who like that good best among available goods.
- ▶ The formal proof makes this intuition precise.

Ordinal inefficiency in RP

Proposition:

In replica economies, if RP is ordinally efficient/inefficient in the base economy (i.e. $q = 1$), then RP is ordinally efficient/inefficient for all replicas.

- ▶ Inefficiency of RP does not disappear completely in any finite replica economy, if RP is inefficient in the base economy.
- ▶ But the “magnitude” of ordinal inefficiency vanishes as markets become large.
- ▶ Manea (2008): Probability that RP fails exact ordinal efficiency goes to zero as the market becomes large.

Extensions

Existing Priorities (e.g., university house allocation with priority for freshmen)

- ▶ Asymmetric RP: agents draw numbers from different distributions, reflecting the priority structure.
- ▶ Asymmetric PS: agents have different eating speeds.
- ▶ The asymmetric RP and the asymmetric PS converge to the same limit.

Multiunit demand: Each agent consumes up to k units (Kojima, 2008).

- ▶ Once-and-for-all RP: Draw a random order. An agent claims k units at her turn.
- ▶ Draft RP: Draw a random order. Each agent claims one at a turn. After all have taken their turns, draw another random order, and so on.
- ▶ Introduce two generalizations of PS, and the two versions of RP above converge to these two versions of PS in the limit.