Designing Random Assignment Mechanisms

Fuhito Kojima[1]

April 16, 2014

[1]Stanford University

Overview

- Social planner wants to assign goods to its members.
  - Student placement mechanisms (NYC, Boston).
  - House allocation in colleges.
- Constraints:
  - No monetary transfer, random assignments.
  - Only ordinal preferences (ranking over goods) can be used.
- There are two competing mechanisms: random priority (RP) and probabilistic serial (PS). We'll study what the issues are and how to evaluate the tradeoffs between them.

Model

- There is a fixed set of **types of goods** O, plus the null good (receiving no good) $\o$.
- Each agent i has strict preferences over $O \cup \{\o\}$.
- A **random assignment** is a matrix $P = [P_{ia}]_{i,a}$ where P_{ia} is the probability that agent i obtains good a.

Random Priority Mechanism

- **Random Priority (RP)**:
  1. Draw each ordering of the agents with equal probability, and
  2. The first agent receives her most preferred good, the next agent his most preferred good among the remaining ones, and so on.
- RP is
  - Easy to implement.
  - Strategy-proof.
  - Fair (equal treatment of equals).
  - Ex post efficient.
  - Widely used in practice.

RP is Ex Ante Inefficient

- Real goods $O = \{a, b\}$ with one copy each and agents $N = \{1, 2, 3, 4\}$,
  1 and 2 like $a, b, \emptyset$ (in this order),
  3 and 4 like $b, a, \emptyset$.
- The random assignments under *RP*

	Good a	**Good** b	**Good** $\emptyset$
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** $\emptyset$
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,

$$\begin{aligned} & \Pr[i \text{ gets } a \text{ or more preferred under } P] \\ \geq \; & \Pr[i \text{ gets } a \text{ or more preferred under } P'], \quad \text{for all } i, a, \end{aligned}$$

  with strict inequality for at least one pair i, a.
- A random assignment is **ordinally efficient** (sd-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 (last example).

Probabilistic Serial Mechanism

- Bogomolnaia and Moulin (2001) define PS based on an "eating algorithm":
  1. Imagine each good is a divisible good of "probability shares."
  2. Imagine there is a time interval $[0, 1]$.
  3. Each agent "eats" the best good with speed one at every time (among goods that have not been completely eaten away).
  4. At time $t = 1$, each agent is endowed with probability shares.
  5. PS assignment is the resulting profile of shares.

PS Mechanism Example

- The same example as before: $O = \{a, b\}$, with one copy each, $N = \{1, 2, 3, 4\}$,
  1 and 2 like $a, b, \emptyset$ (in this order),
  3 and 4 like $b, a, \emptyset$.
- Compute the PS assignment:
  1. $t = 0$: Agents 1 and 2 start eating a, and agents 3 and 4 start eating b.
  2. $t = 1/2$: goods a and b are eaten away. No (real) goods remain.
- The resulting assignment

	Good a	**Good** b	**Good** $\emptyset$
Agents 1 and 2	1/2	0	1/2
Agents 3 and 4	0	1/2	1/2

  is ordinally efficient.

PS is Ordinally Efficient

More generally,

Theorem

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

So PS eliminates (ordinal) inefficiency, unlike RP.

Proof sketch

An intuitive reason that PS is ordinally efficient comes from the nature of "eating."

At each moment of time in the time interval, everyone is eating (infinitesimal share of) his or her favorite available good.

Suppose that agent i likes a better than b but eats b. Then a was already eaten away, and so on.

PS is fair

We can think of an appropriate fairness criterion for random assignments.

We say that a random assignment is **envy-free** if everyone likes his or her assignment better than assignment of anyone else (in the sense that her assignment stochastically dominates others').

Theorem

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

Proof sketch

An intuitive reason that PS is envy-free is, again, from the nature of "eating."

At each moment of time in the time interval, everyone is eating (infinitesimal share of) his or her favorite available good.

So, everyone has chance to eat a better (from his/her viewpoint) good than anyone else, so at the end, no one envies assignments of someone else.

PS is not strategy-proof

- goods $\{a, b, \o\}$ with one unit each, $N = \{1, 2, 3, 4\}$,
  1 likes $a, b, \o$ (in this order),
  2 likes $a, \o$,
  3,4 like $b, \o$.
- If 1 reports true preferences,

	Good a	**Good** b	**Good** $\o$
Agents 1 and 2	1/2	0	1/2
Agents 3 and 4	0	1/2	1/2

- If 1 reports a lie: $b, a, \o$,

	Good a	**Good** b	**Good** $\o$
Agent 1	1/3	1/3	1/3
Agent 2	2/3	0	1/3
Agents 3 and 4	0	1/3	2/3

1 may be made better off.

Impossibility Result

So, PS is better in efficiency (ordinal efficiency), but RP is better in incentives (strategy-proofness).

Is there any other mechanism that can dominate both RP and PS?

Theorem (Bogomolnaia and Moulin)

There is no mechanism that satisfies ordinal efficiency, strategy-proofness and equal treatment of equals.

So, in a sense, we cannot avoid the tradeoff by trying to find another mehcanism.

Then which mechanism, if any, should we use?

NYC's data

Pathak (2008) used NYC's data to compare RP and PS:

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

Pathak notes the difference 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 it, he supported RP over PS.

PS becomes strategy-proof in large markets

Kojima and Manea show that PS becomes strategy-proof in large markets (motivated by applications like NYC)

Theorem (Kojima and Manea 2010)

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 a finitely large markets.

The bound M can be reasonably small: 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 of the Theorem

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.

Asymptotic Equivalence

Consider a sequence of economies, where a **q-economy** is composed of

- q copies of each (real) good and infinite copies of ø, and
- Set of agents: we only assume

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

converges as $q \to \infty$ 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 $q \to \infty$.
Formally,
$\lim_{q\to\infty} \max_{\pi,a} |RP_a^q(\pi) - PS_a^q(\pi)| = 0$, *where*

$$RP_a^q(\pi) := \Pr[\textit{agents with preference } \pi \textit{ get } a \textit{ in } q\textit{-economy in } RP],$$
$$PS_a^q(\pi) := \Pr[\textit{agents with preference } \pi \textit{ get } a \textit{ in } q\textit{-economy in } PS].$$

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

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

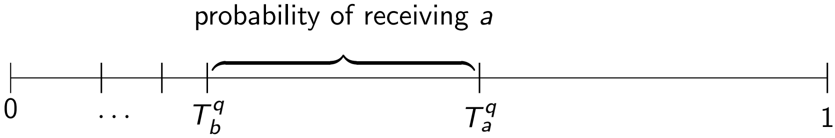

Intuition of the Theorem (2)

- 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 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

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

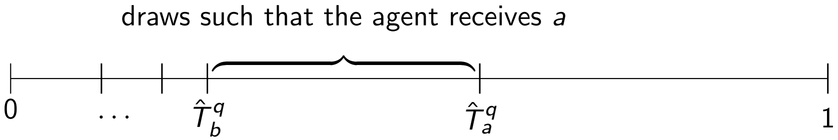

Intuition of the Theorem (3)

- We show 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.

Nature of Ordinal Inefficiency in RP

- We need to make a convergence argument, because for any finite size, RP and PS may not be exactly equivalent.
- Consider a family of replica economies (i.e. agents of each preference type increase proportionately to q).

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.

- Thus, inefficiency of RP does not disappear completely in any finite replica economy, if RP is inefficient in the base economy.
- But our theorem says that the "magnitude" of ordinal inefficiency vanishes as markets become large.
- Manea (2008): Probability that RP fails exact ordinal efficiency goes to one as the market becomes large.

Example

- Consider replica economies of the previous example: $\tilde{O} = \{a, b, \emptyset\}$. The probability of obtaining less preferred good is positive for all q but approaches zero as $q \to \infty$.

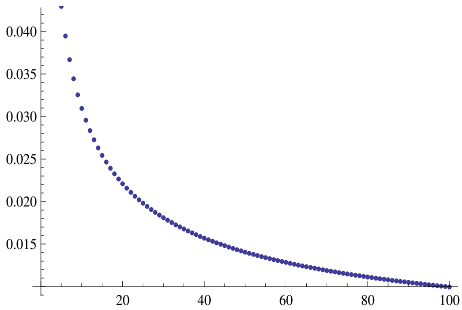

Figure : Horizontal axis: Market size q. Vertical axis: $RP_b^q(\pi)$.

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 $k \geq 2$ 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.
  - We can introduce two generalizations of PS, and the two versions of RP above converge to these two versions of PS in the limit.

Implementing Random Assignments

- A **random assignment** is a matrix $P = (P_{ia})$ where P_{ia} represents the probability that agent i receives good a (i.e., giving marginal distributions).
- It's often easier to work with random assignments directly rather than lotteries over sure outcomes, when considering (utilitarian) welfare, fairness, incentives, etc.
- A question: is dealing with random assignments "justified?" I.e., given a random assignment, is there always a lottery over sure outcomes that realizes it?

Example: A Simple Environment

- Consider a problem where each of n agents receives exactly one of n goods each.
- Each deterministic assignment can be expressed by a **permutation matrix**, i.e., a matrix where
  - Each entry is 0 or 1 (1 means receiving a good).
  - Each row sums up to one (each agent must receive exactly one good).
  - Each column sums up to one (each object must be assigned to exactly one agent).
- Any random assignment in this problem is a convex combination of permutation matrices, so it should be a **bistochastic matrix**, i.e.,
  - Each entry is nonnegative.
  - Each row sums up to one.
  - Each column sums up to one.

Birkhoff-von Neumann Theorem

Theorem (Birkhoff-von Neumann)

Every bistochastic matrix can be written as a convex combination of permutation matrices.

- This theorem implies that any random assignment can be implemented (why)?
- The theorem implies that any random assignment is induced by a lottery assignment.

Proof

Lemma

Let P be a bistochastic matrix that is not a permutation matrix. Then it can be written as a convex combination of two bistochastic matrices,

$$P = \lambda P^1 + (1 - \lambda)P^2,$$

where P^1 and P^2 has the following properties:

1. *If P_{ia} is an integer, then P^1_{ia} and P^2_{ia} are integers.*
2. *P^1 and P^2 has at least one more integral entry than P.*

- Note that the lemma implies the theorem.
- Also, the proof of the lemma provides a practical method to execute the lottery.
- Proof: Blackboard (see Hylland and Zeckhauser 1979 for example)

What Random Assignments Can Be Implemented in General?

Based on Budish, Che, Kojima and Milgrom (2009).

- In many applications, there are more complicated constraints.
  - Many-to-one assignment: Multiple seats in each school.
    $\Rightarrow$ Constraint may be an integer different from one
  - Non-assignment: Opt out to private school.
    $\Rightarrow$ Constraint may be an inequality, not equality
  - Group-specific quota ("Controlled choice"): Affirmative action, Gender Balance, Test score balance, District Favoritism
    $\Rightarrow$ Sub-column constraint
  - Flexible capacity: the relative sizes of alternative programs across schools or within each school may be adjustible.
    $\Rightarrow$ Multi-column constraint
- Budish et al. identify a condition on the constraint structure (called "bihierarchy") that is necessary and sufficient for the existence of a lottery over sure outcomes that implements any given random assignment.

Other topics

- Extending the PS mechanism:
  - Weak preferences: Katta and Sethuraman (2006)
  - Initial endowment ("existing tenants"): Yilmaz (2008, 2009)
  - Group-specific quotas, flexible production, etc: Budish, Che, Kojima and Milgrom (2010), Ball (in progress)
  - Priorities: Kesten and Unver (2009)
- Characterizing the PS mechanism (Hashimoto and Hirata, Kesten, Kurino, and Unver 2010, Bogomolnaia and Heo 2010)
- An equivalent representation: Kesten (2006),
- Characterizations of ordinal efficiency: McLennan (2002), Manea (2006), Carroll (2013), Abdulkadiroglu and Sonmez (2003), Ball (2013)
- Stronger efficiency: Featherstone (2013)
- Incentives: Balbuzanov (2014), Mennle and Seuken (2014a, 2014b)
- Implementing assignment; Akbarpour and Nikzad (2014), Pycia and Unver (2014)

Summary

- The popular RP mechanism is challenged by the new PS mechanism, which may be more efficient.
- PS becomes strategy-proof in large markets, and RP and PS mechanisms converge to the same limit as the copies of each good becomes infinitely large.
- Both RP and PS may be good solutions in large markets.
- Further topics:
  - How well these mechanisms work in large finite economies.
  - Asymptotic equivalence for other mechanisms.
  - Consider more practical (more "complex") environments, for applications?