Indivisible objects allocation

Market Design Mini-Course, HUJI
Assaf Romm

Discrete objects allocation

- Differences: no priorities, multiple objects, different goals

- Extending our toolbox: Top Trading Cycles, (Random) Serial Dictatorship, Competitive equilibrium mechanism, Probabilistic Serial and ordinal efficiency, Rank efficiency.

- Equivalences and large markets results.

Some applications to keep in mind

- House allocation (sounds ridiculous, but think dorms)
- School choice
- Kidney exchanges
- Teach for America
- HBS FIELD 3
- Medical internships in Israel

- Course allocation in some business schools

- And more…

Top Trading Cycles

- Scarf and Shapley (1974)
- We already discussed this, so just a quick reminder:
  - Objects are somehow initially allocated
  - Then each agent points to most preferred object, and cycles are removed (think of this as trades among the top cycles…).

(Random) Serial Dictatorship

- Easiest algorithm ever
- Pick some order over agents (perhaps, randomly).
- Each agent picks an object in turn from the pool of objects left.
- Obviously strategyproof

- RSD is also known as Random Priority (RP)

Theorem (Abdulkadiroglu and Sonmez, 1998): RSD and TTC with random endowments are equivalent.

Competitive market equilibrium

- Based on Hylland and Zeckhauser (1979)
- To implement all the algorithms we discussed up until now (DA, TTC, RSD, Boston), all that was needed was reports of ordinal preferences.
- However, if cardinal preferences can be provided, a better allocation can be achieved.
- Suppose utility functions were reported by participants, and now create a market for probability shares on objects. Each person is given a budget, and can buy (divisible) probability shares. We find a price vector that clears the market. Then we perform a lottery (possible by Birkhoff-von Neumann theorem).
- Problems: Not strategyproof, requires very detailed report

Ordinal efficiency

- Based on Bogomolnaia and Moulin (2001).
- The idea is to get as much as we can from ordinal preferences, in a similar way to Hylland and Zeckhauser.
- What does it mean as much as we can?
- Suppose there are four objects $\{a, b, c, d\}$ and four agents $\{1,2,3,4\}$ with the following preferences:

Agents 1,2: $a \succ b \succ c \succ d$

Agents 3,4: $b \succ a \succ d \succ c$

  - What would RSD give?
  - What would be a better allocation?

<u>Ordinal efficiency</u> – the resulting distribution is not stochastically dominated.

Probabilistic Serial (PS)

- The solution: "simultaneous eating algorithms"
  - Set a vector of cumulative probabilities left of each object, initialized it with ones.
  - Start at time 0 and in continuous time let each participant eat from his/her most preferred object's cumulative probability (as long as it is not zero).
  - For standard PS – select equal eating rates.
  - When all items have been eaten (or when each participant ate measure 1 of objects) – take the resulting matrix, apply the Birkhoff-von Neumann theorem, and draw an allocation.

Claim: The resulting distribution is ordinally-efficient. For standard PS, the resulting distribution is (ex-ante) envy-free.

Probabilistic Serial (PS)

<u>Claim</u>: The PS mechanism is not strategyproof (but it is weakly strategyproof, meaning, you cannot get a distribution that stochastically dominates truthful revelation).

<u>Partial proof</u>: Consider the following example:

$$1{,}2\colon a \succ b \succ c$$
$$3 : b \succ c \succ a$$

If all agents report truthfully agent 1 gets $\frac{1}{2}a + \frac{1}{6}b + \frac{1}{3}c$. If he reports $b \succ a \succ c$, agent 1 gets $\frac{1}{2}b + \frac{1}{4}a + \frac{1}{4}c$. For certain vNM utility functions, agent 1 is better off after the manipulation (e.g. $u_1(a) = 10, u_1(b) = 9, u_1(c) = 0$).

Large market results

- Kojima and Manea (2008): PS is (nearly) strategyproof for large enough finite markets.
- Manea (2009): When there are many object types, the fraction of preferences profiles for which RSD is ordinally efficient tends to zero. Similarly, when the number of objects is constant and there are many copies of each object.
- Che and Kojima (2010): In large markets (such as the ones analyzed by Manea, 2009) the allocations of RSD and PS converge.
- Liu and Pycia (2012): If two mechanisms are efficient, symmetric and asymptotically strategyproof, then they coincide asymptotically.

Rank efficiency

- Based on Featherstone (2011).
- Some organizations really care a lot about the rank distribution that the matching mechanism induces.
- The Teach for America experience

Rank efficiency – the resulting rank distribution is not stochastically dominated.

Claim: All rank efficient distribution are ordinal efficient, but not vice-versa.

Intuition: Rank efficiency is a "social welfare" criterion (you are willing to sacrifice one agent's utility for another's).

Rank efficiency

- How to reach a rank-efficient allocation?
  - Method 1: Start with some allocation, and try to preform "trade cycles" (but this time, only the rank distribution should be better, not each agent)
  - Method 2: Solve a linear programming problem

- As with ordinal efficiency, neither method will be strategyproof (in fact, there is a impossibility result).

Rank efficiency

- Results from HBS (900 students, 11 countries)

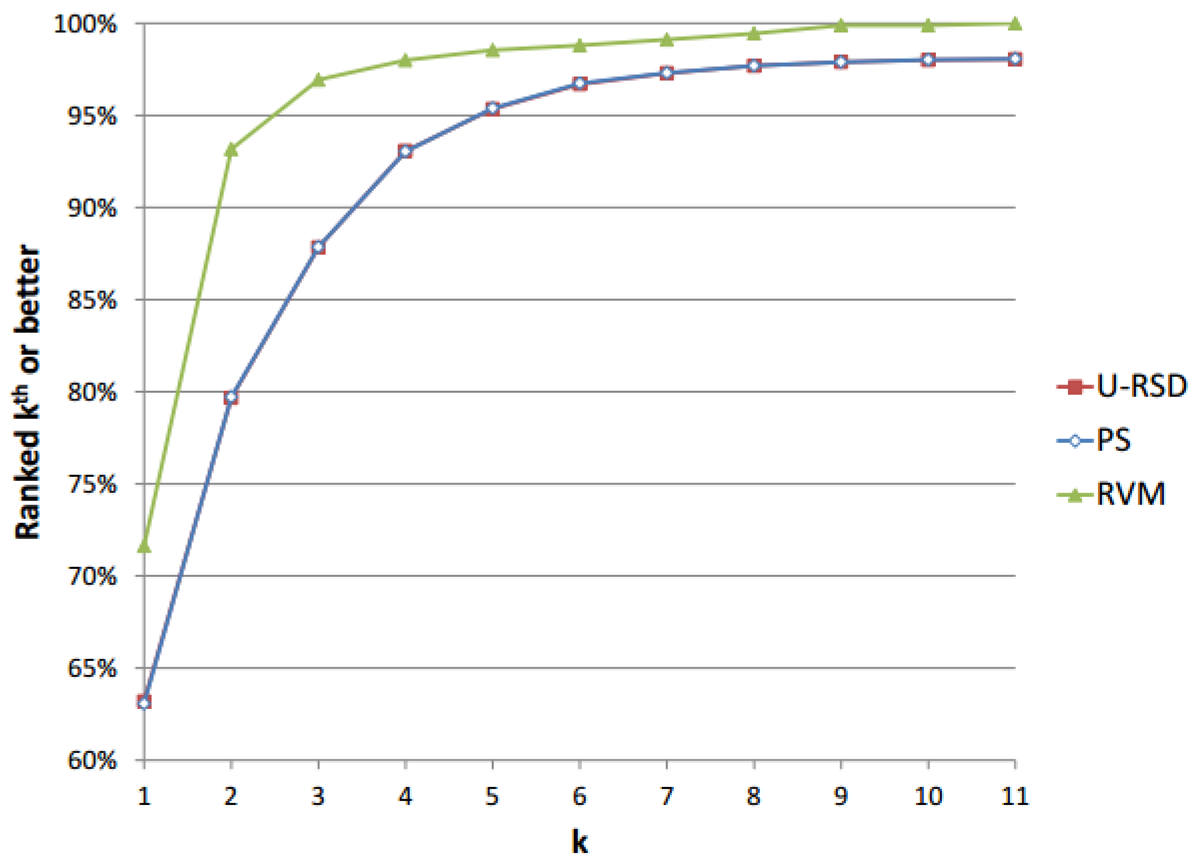

Rank efficiency in large markets

- Do these results persist in large markets?
- Yes!

What did we learn?

- Toolbox
  - DA + Extensions, TTC, RSD, Signaling, Scramble, PS, Rank efficient LP, competitive market mechanism.
- Evaluating solutions
  - Design principles, matching theory, computational experiments, lab experiments, field experiments (including anecdotes and ~~learning from mistakes~~).
- Learning through case studies
  - Medical matches, school choice systems, ~~kidney exchanges~~, and more.