

Allocating Objects

October 26, 2011

Allocating Objects

- ▶ So far we talked about two-sided matching:
 - ▶ Men and Women
 - ▶ Students and Colleges
 - ▶ Workers and Firms
- ▶ When we discuss allocations students to schools: Do schools have preferences?
- ▶ Many public school systems: Schools do not, per se, prefer one child over another. School seats are sometimes viewed as objects to be allocated.
- ▶ House allocation problem.

House allocation problem

A house allocation problem is a tuple (A, H, \succ) where

- ▶ A is a set of agents
- ▶ H is a set of goods (houses) with $|H| = |A|$
- ▶ Each agent $a \in A$ has strict preferences over houses: \succ_a (weak preferences: \succeq_a). $\succ = (\succ_a)_{a \in A}$ is the preference profile.

Possible applications:

- ▶ On campus housing
- ▶ Organ allocation
- ▶ Office allocation
- ▶ School choice problems

Matching μ is a function specifying who receives what house: $\mu(a)$ is the house agent a receives in matching μ .

A matching is Pareto-efficient if: \nexists no other matching ν such that

1. $\nu(a) \succeq_a \mu(a)$ for every agent $a \in A$
2. $\exists a \in A$ such that $\nu(a) \succ_a \mu(a)$

If such an alternative matching ν exists, then ν Pareto dominates μ

Mechanisms

A deterministic mechanism is a rule that for every preference profile \succ assigns a matching. $\varphi(\succ)$ is the matching under mechanism φ when the agents report preference profile \succ .

A mechanism φ is strategy proof if revealing preferences truthfully is a dominant strategy.

A mechanism is Pareto-efficient if $\varphi(\succ)$ is Pareto-efficient for every preference profile \succ .

Serial Dictatorship

A serial dictatorship mechanism (or sometimes priority mechanism) specifies an order over agents, and then lets the first agent receive her favorite good, the next agent receive her favorite good among remaining objects, etc.

The serial dictatorship mechanism is very easy to implement: decide the order (randomly, or using some existing priority such as seniority) and let applicants choose according to the order. Serial dictatorship is used in many applications (with some variations: discussed later): office allocation for professors, NYC school choice system and Columbia and Harvard housing allocation etc.

In addition, serial dictatorship has several good properties.

Theorem: Serial Dictatorship is strategy-proof

Proof: by contradiction:

Suppose there is matching ν that Pareto-dominates μ . Consider the agent a with the highest priority who receives a strictly better object under ν than under μ .

It has to be that

1. There exists an agent b who receives $\nu(b) = \mu(b)$ who chooses before a chooses, else a would have picked $\nu(a)$ under μ .
2. $\nu(b) = \mu(b)$: Because $\nu(b) \succeq_b \mu(b)$ by assumption that ν Pareto dominates μ , and it can't be that $\nu(b) \succ_b \mu(b)$ by definition of a .

Contradiction

Theorem: Serial dictatorship is strategy-proof.

Let the priority order be π .

The first agent $\pi(1)$ of the priority order obtains the favorite good for her when she tells the truth, so, she has no incentives to lie.

The second agent $\pi(2)$ of the priority order obtains the favorite good among the remaining goods for her when she tells the truth, so, she has no incentives to lie. etc.

Axiomatic Characterization

Serial dictatorship has good properties such as strategy-proofness and Pareto efficiency. Are there other mechanisms that have those properties?

There is a whole literature on axiomatic characterization of mechanisms: Find a set of properties (axioms) that are necessary and sufficient for the mechanism to be in a certain class of mechanisms.

For example: for serial dictatorship: which axioms would those be?

Turns out: There are other mechanisms that are Pareto-efficient and strategy-proof.

One way: Find other properties of serial dictatorship.

Group Strategy proofness

Consider the possibility that a group of agents colludes and misreports preferences jointly. Can we assure a mechanism to be immune to such joint manipulations?

Let $\succ_B = (\succ_a)_{a \in B}$ and $\succ_{-B} = (\succ_a)_{a \in A \setminus B}$.

A mechanism φ is group strategy-proof if there is no group of agents $B \subset A$ and preferences \succ'_B such that

1. $\varphi(\succ'_B, \succ_{-B}) \succeq_a \varphi(\succ_B, \succ_{-B})$ for all $a \in B$ and
2. $\varphi(\succ'_B, \succ_{-B}) \succ_a \varphi(\succ_B, \succ_{-B})$ for at least one $a \in B$.

That is, a mechanism is group strategy-proof if no group of agents can jointly misreport preferences in a way to make some member strictly better off, while no one in the group is made worse off.

Theorem: Serial dictatorship is group-strategyproof.

Intuition: Serial dictatorship only uses an agents' preference information when it is the turn for the agent to make a choice. So, the best thing the agent can do: report the most desirable remaining object as most desirable of all remaining objects. Then the mechanism proceeds just as if the agent had report preferences truthfully.

Serial dictatorship is neutral, that is, the assignment does not depend on a “label” (or the name) of goods. For example, the first agent $\pi(1)$ obtains her favorite house, independent of relabeling the houses.

Formally: Let r be a permutation of houses, so $r(h)$ is how house h is called under relabeling r . Let \succ^r be the preference profile where each house h is renamed $r(h)$. A mechanism φ is neutral if, for any permutation r and any \succ , $\varphi(\succ^r)(a) = r(\varphi(\succ)(a))$ for all $a \in A$.

Axiomatization of Serial Dictatorship

Theorem (Svensson 1998):

A mechanism is group strategy-proof and neutral if and only if is a serial dictatorship.

One side: we already have shown. Converse: A little complicated.

Problem of axiomatizations: Which are the most "natural" characterizations...

Housing Market

So far we assumed: No agent own any house. What is agents start with an initial allocation of houses?

Shapley and Scarf (1974): Housing market.

A housing market is $((a_k, h_k)_{k=1, \dots, n}, \succ)$ such that

1. $\{a_1, \dots, a_n\}$ is a set of agents and $\{h_1, \dots, h_n\}$ is a set of houses, where agent a_k owns house h_k .
2. Each agent a has strict preferences \succ_a over houses.

A matching μ is a function specifying who gets what good: $\mu(a)$ is the house that agent a receives in μ .

Just as before, define a mechanism, a mechanism that is strategy-proof and a mechanism that is Pareto-efficient.

A matching μ is in the core if there is no coalition of agents B and a matching ν such that

1. For any $a \in B$, $\nu(a)$ is the initial house of some $b \in B$ and
2. $\nu(a) \succeq_a \mu(a)$ for all $a \in B$ and $\nu(a) \succ_a \mu(a)$ for some $a \in B$.

A matching is individually rational if every agent obtains a house that is at least as good as her initial house.

It is immediate to see that

1. Any core matching is individually rational (consider a one-person coalition $B = \{a\}$).
2. Any core matching is Pareto efficient (consider $B = A$).

Theorem: Shapley and Scarf (1974)

There exists a core matching for any housing market.

This is, in a way, the equivalent to the existence theorem by Gale and Shapley (1962) for two-sided matching, just as the core is to stability (which equals the core in many-to-one two-sided matching).

Gale's Top Trading Cycles TTC algorithm

Step 1: Each agent points to his/her first choice house and each house points to its initial owner. There exists at least one cycle and no cycles intersect. Remove all the cycles and assign each agent in a cycle the house he or she is pointing to.

Step t : Each agent points to his/her first choice house among the remaining ones and each house points to its initial owner. There exists at least one cycle and no cycles intersect. Remove all the cycles and assign each agent in a cycle the house he or she is pointing to.

Since everything is finite, the algorithm ends eventually.

Proof of the theorem

Let μ be the resulting matching from TTC. Suppose there is a coalition B that deviates profitably by inducing matching ν .

Consider the subset of agents in B who strictly prefer their allocation under ν to the one in μ , and let a be an agent who is matched first among this subset in the TTC algorithm.

Then $\nu(a)$ is owned by an agent $b \in B$ who is removed by the TTC algorithm in a strictly earlier step (say cycle C_m).

Then, b obtains a house of $b' \in B \cap C_m$ both in ν and μ , ..., $b^k \in B \cap C_m$ obtains $\nu(a)$ both at ν and μ .

Contradiction

Uniqueness of the core matching

Is there any other matching in the core?

Theorem: Roth and Postlewaite 1977

The matching produced by Gale's TTC algorithm is the unique core matching.

Proof: We have already seen that the TTC algorithm finds a core matching, so we will show there is no other core matching.

Consider an arbitrary matching $\nu \neq \mu$, and fix a to be one of the first agents with $\nu(a) \neq \mu(a)$ (according to the order of being matched in TTC).

Let C_m be the set of agents that form a cycle that includes a . Then, any agents b who are matched before C_m satisfy $\nu(b) = \mu(b)$.

By construction of TTC, $\mu(b) \succeq_b \nu(b)$ for all $b \in C_m$ (because, in ν , all preferred goods are allocated to those who are matched before C_m).

Furthermore, since $\nu(a) \neq \mu(a)$ and $a \in C_m$, we have $\mu(a) \succ_a \nu(a)$.

Since, for any $b \in C_m$, $\mu(b)$ is an initial house owned by some other agent in C_m , these facts imply that C_m can profitably deviate from ν by μ .

Properties of TTC

Theorem: Roth (1982)

The TTC algorithm is strategy-proof.

Theorem Ma(1994)

A mechanism is strategy-proof, Pareto-efficient and individually rational if and only if it is TTC.

What if we combine the two housing models: There are some houses that are owned, and some houses that are not, some agents who own houses, and some who do not. Think about dorm allocation where some students are freshmen, and have no house yet, while others are existing tenants.

Abdulkadiroglu and Sonmez (1999): house allocation with existing tenants.

Each agent has strict preferences over houses (and prefers to be matched rather than unmatched)

Random Serial Dictatorship with Squatting Rights.

Used in undergrad housing in many universities.

1. Each existing tenant decides whether they want to participate in the housing lottery or keep the current house. Those who decide to keep their houses are assigned the current houses. All other houses become available for assignment in later steps.
2. An ordering of agents is decided. The ordering may be uniformly random or may favor some subgroup of agents (for example, seniors over juniors).
3. Serial dictatorship is applied to all available houses and agents (except for existing tenants already assigned their current houses).

What do you think of that mechanism?

Problem:

Existing tenants are not guaranteed to get at least as good a house as their current house: Individually irrational!

Some existing tenants may not want to enter the lottery even if they want to move.

This may result in loss of gains from trade, and the resulting matching may not be Pareto efficient.

Some good properties we want for house allocation mechanisms:

- ▶ Pareto efficiency
- ▶ Strategy-proofness
- ▶ Individual rationality

Also, recall that there were good mechanisms in special cases:

- ▶ Serial dictatorship (house allocation problem)
- ▶ Gale's top trading cycles (housing markets)

Can we find mechanisms with the above good properties? Can SD and TTC help?

The You Request My House - I Get Your Turn mechanism

1. Let the agent with the top priority receive her first choice good, the second agent his top choice among the remaining goods and so on, until someone requests the house of an existing tenant.
2. If the existing tenant whose house is requested has already received a house, then proceed the assignment to the next agent. Otherwise, insert the existing tenant at the top of the priority order and proceed with the procedure.
3. If at any step a cycle forms, the cycle is formed by existing tenants $(a_1, ..a_k)$ where a_1 points to the house of agent a_2 , who points to the house of a_3 , and so on. In such a case assign these houses by letting them exchange, and then proceed with the algorithm.

The YRMH-IGYT mechanism generalizes previous mechanisms:

1. Serial dictatorship when there are no existing tenants:
Without existing tenants, the “you request my house...” contingency simply does not happen, so the mechanism coincides with serial dictatorship straightforwardly.
2. Gale’s TTC if all agents are existing tenants and there is no vacant house: In that case, an agent’s request always points to a house owned by someone, and the assignment of a house happens if and only if there is a cycle made of existing tenants.
3. Indeed, we can think of YRMH-IGYT as a variant of Gale’s TTC in which all vacant houses (and houses whose initial owners are already assigned houses) point to the highest priority agents rather than the owners of the houses. So we sometimes call the mechanism TTC as well.

Theorem:

Any TTC mechanism is individually rational, strategy proof and Pareto-efficient.

Proof sketch: As TTC (YGMH-IGYT) is a common generalization of serial dictatorship and Gale's TTC, Pareto efficiency and strategy-proofness are inherited from these mechanisms (the proof is quite similar).

Also, individual rationality is inherited from Gale's TTC, and can be understood as follows: Whenever some agent points to a house of an existing tenant, she is promoted to the top of the priority. Whenever her top choice at this stage is her own house she can keep it by forming a "cycle" composed of herself and her house.