Harvard CS 1360, Spring 2026
Instructor: Ariel Procaccia
Scribe: Kelsey Chen, Jakob de Raaij

Random Assignment

—

Lecture 16

In contrast to the two-sided matchings from last lecture — where both 'sides' of the players to match (i.e., students and courses) had preferences — we now consider one-sided matchings where players have preferences over goods. There is the same number of goods and players and each player receives exactly one good; thus, this is a more restricted version of fair division of indivisible goods, as seen earlier. However, since every player receives exactly one good, envy freeness up to one good (our sought-after fairness notion in the more general setting) is trivial — so can we do better?

This problem is inspired by real-world applications. When allocating public housing to applicants, the housing options have no preferences over potential tenants, while the potential tenants do have preferences over the housing options. In school choice, some students may remain unmatched in earlier allocation rounds (e.g., with a variant of the deferred acceptance algorithm). To distribute these unmatched students to schools, usually only the students' preferences are taken into account.

1 One-Sided Matchings and Serial Dictatorship

Definition 1 (One-Sided Matching)**.** An instance of *one-sided matching* consists of

- a set $N = \{1, ..., n\}$ of n players,
- a set G of n goods (note, the number of players and goods is the same), and
- for each player $i \in N$, a ranking $\sigma_i \in \mathcal{L}$ over the goods G.

An *assignment* is a perfect matching $\pi : N \to G$ between players and goods, where $\pi(i)$ is the good assigned to player i. A *rule (or mechanism)* f takes as input the rankings of each player, $\boldsymbol{\sigma} \in \mathcal{L}^n$, and outputs an assignment π.

The first mechanism that we will consider is *serial dictatorship*, sometimes also called the *priority mechanism*:

Algorithm 1 Serial Dictatorship for τ

Require: τ is an order of the players.

1: In order of τ, the players select their favorite good that is still available.

Example 1 (Serial Dictatorship)**.** Consider the preferences of $n = 4$ players over goods $G = \{a, b, c, d\}$ on the left and the order $1 \succ_\tau 2 \succ_\tau 3 \succ_\tau 4$. According to τ, player 1 gets to choose first and picks their favorite good a. Then, the next player, player 2, can no longer choose their favorite good a, so they pick their next favorite good b. Next, player 3 chooses their favorite still available good d, then player 4 chooses their favorite still available good c. This is shown on the right.

1	2	3	4
a	a	d	a
b	b	c	d
c	c	b	c
d	d	a	b

1	2	3	4
a	~~a~~	d	~~a~~
b	b	c	~~d~~
c	c	~~b~~	c
d	d	~~a~~	~~b~~

We now formally define two desirable properties — that we already know well from other settings — for allocation rules in one-sided matching and show that serial dictatorship satisfies both of those properties.

Definition 2 (Pareto Efficiency)**.** An assignment π is *Pareto efficient* if there is no assignment π' such that $\pi'(i) \succeq_{\sigma_i} \pi(i)$ for all $i \in N$ and $\pi'(j) \succ_{\sigma_j} \pi(j)$ for some $j \in N$. In other words, there exists no other assignment π' that makes no player worse off but some player strictly better off.

A rule f is *Pareto efficient* if it returns a Pareto efficient assignment for all possible input rankings $\boldsymbol{\sigma} \in \mathcal{L}^n$.

Theorem 1. *Serial dictatorship is Pareto efficient for every order τ.*

Proof. Let π be the assignment found by serial dictatorship. Assume towards a contradiction that there exists some other assignment π' that *Pareto dominates* π, i.e., $\pi'(i) \succeq_{\sigma_i} \pi(i)$ for all $i \in N$ and $\pi'(j) \succ_{\sigma_j} \pi(j)$ for some $j \in N$. Let i be the first player in the order τ that gets a different good in π' than in π. Since π' Pareto dominates π, we know that i strictly prefers their good in π' over their good in π. However, since all players before i received the same good in π and π', we know that $\pi'(i)$ was still available in serial dictatorship when i chose $\pi(i)$. This is a contradiction to i choosing their favorite good, so we conclude that no assignment π' can Pareto dominate the assignment output by serial dictatorship. It is Pareto efficient. □

Definition 3 (Strategyproofness)**.** A rule f is *strategyproof* if for all $\boldsymbol{\sigma} \in \mathcal{L}^n$, $i \in N$, and $\sigma'_i \in \mathcal{L}$, it holds that

$$f(\boldsymbol{\sigma})(i) \succeq_{\sigma_i} f(\sigma'_i, \boldsymbol{\sigma}_{-i})(i).$$

In other words, for all input rankings $\boldsymbol{\sigma} \in \mathcal{L}^n$, no player i can receive a strictly preferred good if they misreport their ranking.

Theorem 2. *Serial dictatorship is strategyproof for every order τ.*

Proof. Each player receives their favorite good among the goods still available when it is their turn. Their preferences do not impact the choices of the players before them in τ, so they cannot change the set of goods available to them. Thus, they cannot do better than receiving their favorite good available to them, which is the case if they report their ranking truthfully. □

Thus, serial dictatorship satisfies two desirable properties. However, it feels intuitively unfair since we have to pick an order τ, and players that are earlier in the order τ will have an advantage over players that come later in τ. To avoid this unfairness, we can choose the order in which the players get to pick at random:

Algorithm 2 Random Serial Dictatorship

Pick an order τ of the players uniformly at random.
Allocate the goods according to serial dictatorship with order τ.

Example 2 (Random Serial Dictatorship)**.** Consider the following preferences of $n = 3$ players over goods $G = \{a, b, c\}$

1	2	3
a	a	b
b	b	c
c	c	a

There are six possible orders that lead to the following assignments, each of which is returned with probability $1/6$.

1	2	3
a	a	b
b	b	c
c	c	a

$1 \succ_\tau 2 \succ_\tau 3$

1	3	2
a	b	a
b	c	b
c	a	c

$1 \succ_\tau 3 \succ_\tau 2$

2	1	3
a	a	b
b	b	c
c	c	a

$2 \succ_\tau 1 \succ_\tau 3$

2	3	1
a	b	a
b	c	b
c	a	c

$2 \succ_\tau 3 \succ_\tau 1$

3	1	2
b	a	a
c	b	b
a	c	c

$3 \succ_\tau 1 \succ_\tau 2$

3	2	1
b	a	a
c	b	b
a	c	c

$3 \succ_\tau 2 \succ_\tau 1$

Let's examine whether random serial dictatorship still satisfies the desirable properties discussed. Importantly, however, since now the outcome is random, we need to slightly change the definitions of our properties. In particular, we will say that a rule or mechanism satisfies a property *ex post* if it satisfies this property for any possible realization of the randomness within the rule.[1] In the case of random serial dictatorship, this means that the returned assignment needs to satisfy a property for every possible τ so that it satisfies this property ex post. Since serial dictatorship is strategyproof and Pareto-efficient for every τ, we get the following:

Theorem 3. *Random serial dictatorship is ex post strategyproof and ex post Pareto efficient.*

In other words, no matter which order τ gets chosen at random, no player would have been better off misreporting their ranking.

Definition 4 (Random Assignment)**.** We call a distribution over possible assignments (as is induced by random serial dictatorship) a *lottery*. For any lottery, we define the corresponding *random assignment* to be the matrix $P = [p_{i,x}] \in [0,1]^{n\times n}$ where $p_{i,x}$ is the probability of player i being assigned good x under this lottery. This matrix is *bistochastic*, meaning that the entries in any row and the entries in any column sum to 1.

Example 3 (Random Assignment)**.** Let's consider the lottery induced by random serial dictatorship on the preferences shown in Example 2. This lottery gives the random assignment shown below. For example, we can check that player 1 gets a in three of the six orders, gets b in one of the six orders, and gets c in two of the six orders, leading to the probabilities in the random assignment.

	a	b	c
1	1/2	1/6	1/3
2	1/2	1/6	1/3
3	0	2/3	1/3

Using this definition, we can also formalize our intuition that random serial dictatorship does not favor one player over another player with the same ranking.

Definition 5 (Equal Treatment of Equals)**.** A rule f satisfies *equal treatment of equals* if for any players $i, j \in N$ with $\sigma_i = \sigma_j$, it holds that $p_{i,x} = p_{j,x}$ for all goods $x \in G$. In other words, any two players with the same ranking have the same random allocation probabilities.

Theorem 4. *Random serial dictatorship satisfies equal treatment of equals.*

This theorem follows from the fact that random serial dictatorship is fully symmetric in the players.

2 Ordinal Efficiency

So far, it looks like random serial dictatorship might have all properties we may desire! However, it turns out that ex post Pareto efficiency is not as strong of an efficiency guarantee as we can hope for.

Example 4 (Ordinal Efficiency)**.** Consider the following preferences of $n = 4$ players over goods $G = \{a, b, c, d\}$. Note that all players have almost the same ranking, just the top two alternatives are switched between 1,2 and 3,4.

1	2	3	4
a	a	b	b
b	b	a	a
c	c	c	c
d	d	d	d

[1] In contrast, you may encounter the term *ex ante* if a rule satisfies a property only 'in expectation' but not necessarily for every possible realization of the randomness.

The random serial dictatorship mechanism leads to the random assignment shown below on the left. Note that player 1 receives b if and only if one of players 3,4 chooses first but player 1 is choosing second (so with probability $1/12$). Similarly, players 2, 3, and 4 receive their second-favorite good with probability $1/12$. Now, the players could make the following deal: If one of those situations arises where a player would choose their second-favorite good, the player instead chooses their third-favorite good. While this hurts the given player for this order (for example, player 1 receiving good c instead of b), it will benefit them a lot for other orders (for example, player 1 receiving a instead of c when player 2 chooses first but player 1 does not choose second). Intuitively, the benefit a player gets from another player following the deal in their favor is greater than the harm they suffer from following the deal themselves. Indeed, one can check that random serial dictatorship with this deal leads to the alternate assignment on the right. The only difference is that the probability of $1/12$ with which each player received their second-favorite choice went towards the probability with which the players receive their favorite choice — so every player is strictly better off!

	a	b	c	d
1	$5/12$	$1/12$	$1/4$	$1/4$
2	$5/12$	$1/12$	$1/4$	$1/4$
3	$1/12$	$5/12$	$1/4$	$1/4$
4	$1/12$	$5/12$	$1/4$	$1/4$

Random serial dictatorship

	a	b	c	d
1	$1/2$	0	$1/4$	$1/4$
2	$1/2$	0	$1/4$	$1/4$
3	0	$1/2$	$1/4$	$1/4$
4	0	$1/2$	$1/4$	$1/4$

Alternate random assignment

Let's formalize what it means for all players to be 'strictly better off' under one random assignment over another.

Definition 6 (Stochastic Dominance)**.** A random assignment P' stochastically dominates P if (and only if) for every player $i \in N$ and every good $x \in G$, it holds that

$$\sum_{y \succeq_{\sigma_i} x} p'_{iy} \geq \sum_{y \succeq_{\sigma_i} x} p_{iy}$$

with at least one inequality being strict. In other words, for any player and any *prefix* of their ranking (i.e., their favorite k goods, for some k), the probability that they received a good from this prefix is no smaller under P' than under P, and strictly greater for some player and prefix. This means that any player will be at least as satisfied with their random allocation under P' as under P, and some player will be strictly better off.

You can check that in Example 4 the alternative random assignment, P', stochastically dominates the random assignment returned by random serial dictatorship. For example, consider player 1 and their prefix given by all alternatives they weakly prefer to a, i.e., their top 1 alternative a. The probability that this player receives an alternative from this prefix strictly increases from $5/12$ under P to $1/2$ under P'. At the same time, for no player and no prefix, the probability of this player receiving that prefix went down.

Definition 7 (Ordinal Efficiency)**.** A random assignment is *ordinally efficient* if it isn't stochastically dominated by any other assignment. A rule is *ordinally efficient* if the induced random assignment is ordinally efficient for all possible input rankings $\boldsymbol{\sigma} \in \mathcal{L}^n$.

Theorem 5. *Random serial dictatorship is not ordinally efficient.*

We have seen an example (and thus proof) of this in Example 4. It is also noteworthy that ordinal efficiency is strictly stronger than ex post Pareto efficiency. To see this, assume towards a contradiction that a rule outputs with some non-zero probability an assignment π that is Pareto dominated by some other assignment π'. If we shifted all probability weight from π to π', we obtain a new random allocation that stochastically dominates the old random allocation (we leave formally verifying this as an exercise). Thus, if a rule is ordinally efficient, it is also ex post Pareto efficient.

Now that we have found a strictly stronger efficiency property, some new questions arise: How do we find a ordinally efficient random assignment? Can any random assignment matrix be realized using a lottery over (deterministic) assignments? And can we do this without having to give up on our other desired properties? We will answer these questions now.

3 Probabilistic Serial Rule

Algorithm 3 Probabilistic Serial Rule

1. Assume there is one unit of each good.
2. Every player 'eats' their favorite good at the same rate.
3. When the one unit of a good is fully eaten up, each player that was eating it moves to their next preferred, not yet fully eaten, good and eats it.
4. The process ends when every good is eaten up. The random assignment probability $p_{i,x}$ of player i and good x is the fraction of the good that player i ate.

Example 5 (Probabilistic Serial Rule)**.** Consider the following preferences of $n = 4$ player $N = \{1, 2, 3, 4\}$ over goods $G = \{a, b, c, d\}$.

1	2	3	4
a	b	b	b
b	c	c	d
c	d	d	c
d	a	a	a

Player 1 starts eating good a while players 2,3,4 start eating good b. After 1/3 unit of time, good b is eaten up:

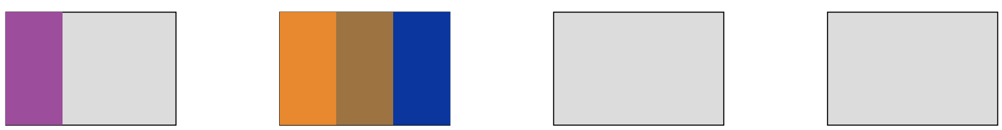

Good a Good b Good c Good d

Players 2 and 3 now go to eat good c, while player 4 goes to eat good d. After 1/2 unit of time, good c is eaten up:

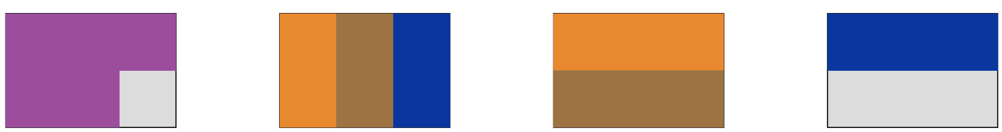

Good a Good b Good c Good d

Players 2 and 3 now join player 4 in eating good d. After 1/6 unit of time, all goods are eaten up:

Good a Good b Good c Good d

This leads to the following random assignment

	a	b	c	d
1	1	0	0	0
2	0	1/3	1/2	1/6
3	0	1/3	1/2	1/6
4	0	1/3	0	2/3

However, to turn this into a rule, we still need to turn the random assignment matrix into a lottery over assignments.

Definition 8 (Permutation Matrix)**.** A permutation matrix is a bistochastic matrix containing only zeros and ones. In other words, it is a matrix that contains exactly one 1 in each row and in each column, and is 0 everywhere else.

Theorem 6 (Birkhoff-von Neumann)**.** *Any bistochastic matrix P can be expressed as a convex combination of permutation matrices $M_1, ..., M_r$. That is, $P = \sum_{i=1}^{r} \lambda_i M_i$ where $\sum_{i=1}^{r} \lambda_i = 1$ and $\lambda_i \geq 0$ for all i.*

Example 6 (Bistochastic Matrix Decomposition)**.** Let's consider a bistochastic matrix that represents a random assignment, shown on the left below. It can be written as a convex combination of three permutation matrices (shown is one of multiple possible ways) as follows:

$$\begin{array}{c|ccc} & a & b & c \\ \hline 1 & 1/2 & 1/6 & 1/3 \\ 2 & 1/2 & 1/2 & 0 \\ 3 & 0 & 1/3 & 2/3 \end{array} = 1/6 \cdot \begin{array}{c|ccc} & a & b & c \\ \hline 1 & 0 & 1 & 0 \\ 2 & 1 & 0 & 0 \\ 3 & 0 & 0 & 1 \end{array} + 1/2 \cdot \begin{array}{c|ccc} & a & b & c \\ \hline 1 & 1 & 0 & 0 \\ 2 & 0 & 1 & 0 \\ 3 & 0 & 0 & 1 \end{array} + 1/3 \cdot \begin{array}{c|ccc} & a & b & c \\ \hline 1 & 0 & 0 & 1 \\ 2 & 1 & 0 & 0 \\ 3 & 0 & 1 & 0 \end{array}.$$

Since each permutation matrix is an assignment, we can go from any random assignment to a lottery over assignments. The coefficients of the assignments in the convex combinations give us the probability of each assignment in this lottery.

We can compute such a decomposition into permutation matrices efficiently, and thus can calculate the probabilistic serial rule efficiently. At a high level, we can repeatedly find a permutation matrix in the bistochastic matrix and then subtract this permutation matrix, multiplied by the smallest corresponding element in the original matrix. In each step, at least one more entry of the bistochastic matrix will equal 0, so after we subtracted at most n^2 permutation matrices, we found a full decomposition of the original matrix into permutation matrices.

The probabilistic serial rule fulfills the desired efficiency guarantee that random serial dictatorship failed at.

Theorem 7. *The probabilistic serial rule is ordinally efficient.*

By symmetry, it is also not hard to check that the probabilistic serial rule does not prefer certain players.

Theorem 8. *The probabilistic serial rule satisfies equal treatment of equals.*

However, it fails another desired property.

Theorem 9. *The probabilistic serial rule is not strategyproof.*

Proof. We prove this by giving an instance where a player benefits from misreporting their true ranking. Consider the following preferences of $n = 4$ players over goods $G = \{a, b, c, d\}$ on the left. The matrix on the right shows the random assignment resulting from the probabilistic serial rule.

1	2	3	4
a	*a*	*b*	*b*
b	*c*	*c*	*c*
c	*d*	*d*	*d*
d	*b*	*a*	*a*

→

	a	*b*	*c*	*d*
1	1/2	0	1/4	1/4
2	1/2	0	1/4	1/4
3	0	1/2	1/4	1/4
4	0	1/2	1/4	1/4

Let's assume that player 1 misreports their preferences by swapping a and b. Intuitively, if they report truthfully, they eat half of a, but by the time they would make it to b it is already fully eaten up. If, instead, they started eating b, they would get 1/3 of b and then still also get 1/3 of a (because now only player 2 is eating a, whereas otherwise players 3 and 4 are both eating b). The preference profile with player 1 misreporting accordingly and the resulting random assignment are shown below.

1	2	3	4
~~a~~ b	a	b	b
~~b~~ a	c	c	c
c	d	d	d
d	b	a	a

$\rightarrow$

	a	b	c	d
1	1/3	1/3	1/12	1/4
2	2/3	0	1/12	1/4
3	0	1/3	5/12	1/4
4	0	1/3	5/12	1/4

In particular, the probability of player 1 being allocated one of a or b (i.e., their prefix of their top 2 goods) increased from 1/2 to 2/3 due to their misreport. Simultaneously, the probability of them being allocated their favorite good a (their prefix of their top 1 good) decreases from 1/2 to 1/3. If player 1 is (almost) indifferent between a and b but is not interested in the other goods at all, this would be a strictly better outcome for them. □

Theorem 10. *There is no rule that satisfies ordinal efficiency, strategyproofness and equal treatment of equals.*

If we accept equal treatment of equals as a non-negotiable, then there is an unavoidable tradeoff between ordinal efficiency and strategyproofness. Random serial dictatorship might lead to outcomes that are not ordinally efficient while the probabilistic serial rule allows for strategic manipulation from the players. So, which should we prefer?

In a 2006 paper by Parag A. Pathak and coauthors, they ran both rules on truthful preference data from 8255 students for school assignments in New York City. They found that the outcomes, measured by the distribution of how many students received their i-th favorite choice, were almost identical under both rules. This empirical result is also supported by theoretical evidence: In a paper from 2010, Yeon-Koo Che and Fuhito Kojima showed that the two rules converge to the same random assignments as the instances grow larger. Given that the results of both mechanisms on truthful data are almost identical, it has been argued that random serial dictatorship should be used in practice to avoid strategic behavior (which, in the case of school assignment, may further disadvantage students from a less privileged background); indeed, random serial dictatorship is currently used in practice while the probabilistic serial rule is not.