Order Symmetry: a new fairness criterion for assignment mechanisms

Mark C. Wilson (`https://markcwilson.site`)
(joint work with Rupert Freeman and Geoffrey Pritchard)

Harvard EconCS Seminar 2021-10-08

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Imagine you are hosting a birthday party for a group of children.
- You have purchased a selection of cheap plastic objects that strangely appeal to this age group, and must give one to each child as a gift.
- Naturally, you want to allocate the objects via a process that is fair to each child, while also guaranteeing that, overall, children are satisfied with their allocation.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Formal model

- Let A be a finite set of n agents and O a finite set of n objects.
- Each agent has a strict linear order over all objects; the collection of all such is the profile.
- The house allocation problem: find a mechanism that for each input profile provides a matching between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Key standard axiomatic properties:
  - Pareto efficiency
  - Anonymity
  - Strategyproofness

Common answer: serial dictatorship

- Fix an order on agents and let them choose in turn their favorite remaining item.
- This mechanism is strategyproof, Pareto efficient and easy to implement.
- However it seems very unfair to those who come later in the choosing order.
- For example, the first agent always gets their top choice, while the last must take whatever is left by all the others.
- This bias is independent of the preferences. For some preference profiles, the last agent does fine, but earlier agents always do at least as well as later ones.

Common answer: serial dictatorship

- Fix an order on agents and let them choose in turn their favorite remaining item.
- This mechanism is strategyproof, Pareto efficient and easy to implement.
- However it seems very unfair to those who come later in the choosing order.
- For example, the first agent always gets their top choice, while the last must take whatever is left by all the others.
- This bias is independent of the preferences. For some preference profiles, the last agent does fine, but earlier agents always do at least as well as later ones.

Common answer: serial dictatorship

- Fix an order on agents and let them choose in turn their favorite remaining item.
- This mechanism is strategyproof, Pareto efficient and easy to implement.
- However it seems very unfair to those who come later in the choosing order.
- For example, the first agent always gets their top choice, while the last must take whatever is left by all the others.
- This bias is independent of the preferences. For some preference profiles, the last agent does fine, but earlier agents always do at least as well as later ones.

Common answer: serial dictatorship

- Fix an order on agents and let them choose in turn their favorite remaining item.
- This mechanism is strategyproof, Pareto efficient and easy to implement.
- However it seems very unfair to those who come later in the choosing order.
- For example, the first agent always gets their top choice, while the last must take whatever is left by all the others.
- This bias is independent of the preferences. For some preference profiles, the last agent does fine, but earlier agents always do at least as well as later ones.

Common answer: serial dictatorship

- Fix an order on agents and let them choose in turn their favorite remaining item.
- This mechanism is strategyproof, Pareto efficient and easy to implement.
- However it seems very unfair to those who come later in the choosing order.
- For example, the first agent always gets their top choice, while the last must take whatever is left by all the others.
- This bias is independent of the preferences. For some preference profiles, the last agent does fine, but earlier agents always do at least as well as later ones.

Randomized mechanism

- We can choose the initial order randomly (obtaining the randomized mechanism **RSD**).
- This method seems fair *ex ante*, but still not fair *ex post*.
- Note that this is a (uniform) lottery over deterministic mechanisms.

Randomized mechanism

- We can choose the initial order randomly (obtaining the randomized mechanism **RSD**).
- This method seems fair *ex ante*, but still not fair *ex post*.
- Note that this is a (uniform) lottery over deterministic mechanisms.

Randomized mechanism

- We can choose the initial order randomly (obtaining the randomized mechanism **RSD**).
- This method seems fair *ex ante*, but still not fair *ex post*.
- Note that this is a (uniform) lottery over deterministic mechanisms.

Fractional mechanism

- Discrete mechanisms have some drawbacks.
- For example, they cannot treat agents symmetrically in all cases (e.g. if preferences are unanimous: all people have the same strict preference order over items).
- A fractional allocation mechanism allocates fractions of objects in a consistent way.
- A fractional allocation can be interpreted as timesharing or dividing objects.

Fractional mechanism

- Discrete mechanisms have some drawbacks.
- For example, they cannot treat agents symmetrically in all cases (e.g. if preferences are unanimous: all people have the same strict preference order over items).
- A fractional allocation mechanism allocates fractions of objects in a consistent way.
- A fractional allocation can be interpreted as timesharing or dividing objects.

Fractional mechanism

- Discrete mechanisms have some drawbacks.
- For example, they cannot treat agents symmetrically in all cases (e.g. if preferences are unanimous: all people have the same strict preference order over items).
- A fractional allocation mechanism allocates fractions of objects in a consistent way.
- A fractional allocation can be interpreted as timesharing or dividing objects.

Fractional mechanism

- Discrete mechanisms have some drawbacks.
- For example, they cannot treat agents symmetrically in all cases (e.g. if preferences are unanimous: all people have the same strict preference order over items).
- A fractional allocation mechanism allocates fractions of objects in a consistent way.
- A fractional allocation can be interpreted as timesharing or dividing objects.

Fractional mechanisms and randomized mechanisms

- Given a randomized mechanism, we can always derive a fractional mechanism (by taking expectations).
- Conversely (Birkhoff-von Neumann theorem): every fractional allocation can be realized as a lottery over discrete allocations.
- A key question is: can we do this in a consistent way so that our fractional mechanism is the expectation of a randomized mechanism where all discrete mechanisms in the support have desirable properties?

Fractional mechanisms and randomized mechanisms

- Given a randomized mechanism, we can always derive a fractional mechanism (by taking expectations).
- Conversely (Birkhoff-von Neumann theorem): every fractional allocation can be realized as a lottery over discrete allocations.
- A key question is: can we do this in a consistent way so that our fractional mechanism is the expectation of a randomized mechanism where all discrete mechanisms in the support have desirable properties?

Fractional mechanisms and randomized mechanisms

- Given a randomized mechanism, we can always derive a fractional mechanism (by taking expectations).
- Conversely (Birkhoff-von Neumann theorem): every fractional allocation can be realized as a lottery over discrete allocations.
- A key question is: can we do this in a consistent way so that our fractional mechanism is the expectation of a randomized mechanism where all discrete mechanisms in the support have desirable properties?

Example

Consider the profile where agents a_1, a_2, a_3 have respective preferences over objects o_1, o_2, o_3 as follows.

$$a_1 : o_1 \succ o_2 \succ o_3$$
$$a_2 : o_1 \succ o_3 \succ o_2$$
$$a_3 : o_2 \succ o_1 \succ o_3$$

Let rows represent agents a_1, a_2, a_3 and columns represent objects o_1, o_2, o_3. The matrix of the fractional assignment from **RSD** is:

$$\begin{bmatrix} 3/6 & 1/6 & 2/6 \\ 3/6 & 0 & 3/6 \\ 0 & 5/6 & 1/6 \end{bmatrix}.$$

Top Trading Cycles (TTC) mechanism

- Each agent is given an initial allocation (the endowment).
- Each agent points to the owner of their favorite item.
- This creates a directed graph which must have at least one cycle.
- Resolve all cycles by giving everyone in a cycle their desired item.
- Continue with the remaining agents, after removing the satisfied ones and their items.

Top Trading Cycles (TTC) mechanism

- Each agent is given an initial allocation (the endowment).
- Each agent points to the owner of their favorite item.
- This creates a directed graph which must have at least one cycle.
- Resolve all cycles by giving everyone in a cycle their desired item.
- Continue with the remaining agents, after removing the satisfied ones and their items.

Top Trading Cycles (TTC) mechanism

- Each agent is given an initial allocation (the endowment).
- Each agent points to the owner of their favorite item.
- This creates a directed graph which must have at least one cycle.
- Resolve all cycles by giving everyone in a cycle their desired item.
- Continue with the remaining agents, after removing the satisfied ones and their items.

Top Trading Cycles (TTC) mechanism

- Each agent is given an initial allocation (the endowment).
- Each agent points to the owner of their favorite item.
- This creates a directed graph which must have at least one cycle.
- Resolve all cycles by giving everyone in a cycle their desired item.
- Continue with the remaining agents, after removing the satisfied ones and their items.

Top Trading Cycles (TTC) mechanism

- Each agent is given an initial allocation (the endowment).
- Each agent points to the owner of their favorite item.
- This creates a directed graph which must have at least one cycle.
- Resolve all cycles by giving everyone in a cycle their desired item.
- Continue with the remaining agents, after removing the satisfied ones and their items.

Properties of Top Trading Cycles mechanism

- Pareto efficient
- Strategyproof
- How fair is it?
- It seems less obviously unfair, since if the endowment is chosen arbitrarily there is no obvious advantage to any agent.
- We can also choose the endowment uniformly at random (**TTC**).

Properties of Top Trading Cycles mechanism

- Pareto efficient
- Strategyproof
- How fair is it?
- It seems less obviously unfair, since if the endowment is chosen arbitrarily there is no obvious advantage to any agent.
- We can also choose the endowment uniformly at random (**TTC**).

Properties of Top Trading Cycles mechanism

- Pareto efficient
- Strategyproof
- How fair is it?
- It seems less obviously unfair, since if the endowment is chosen arbitrarily there is no obvious advantage to any agent.
- We can also choose the endowment uniformly at random (**TTC**).

Properties of Top Trading Cycles mechanism

- Pareto efficient
- Strategyproof
- How fair is it?
- It seems less obviously unfair, since if the endowment is chosen arbitrarily there is no obvious advantage to any agent.
- We can also choose the endowment uniformly at random (**TTC**).

Properties of Top Trading Cycles mechanism

- Pareto efficient
- Strategyproof
- How fair is it?
- It seems less obviously unfair, since if the endowment is chosen arbitrarily there is no obvious advantage to any agent.
- We can also choose the endowment uniformly at random (**TTC**).

The fractional mechanisms are the same

- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante*.
- However we have seen that *ex post* the randomized mechanisms behave differently from the point of view of fairness.
- There is a missing concept!

The fractional mechanisms are the same

- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante*.
- However we have seen that *ex post* the randomized mechanisms behave differently from the point of view of fairness.
- There is a missing concept!

The fractional mechanisms are the same

- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante*.
- However we have seen that *ex post* the randomized mechanisms behave differently from the point of view of fairness.
- There is a missing concept!

The fractional mechanisms are the same

- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante*.
- However we have seen that *ex post* the randomized mechanisms behave differently from the point of view of fairness.
- There is a missing concept!

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Assumptions

- We consider randomized mechanisms where the randomness is realized *before* preferences are revealed.
- Thus the roles of the agents are known before their preferences are. Advantages:
  - lower communication complexity
  - allows agents to more easily reason about their decisions
  - de-emphasizes the randomization
- This decoupling happens in real-world examples, e.g. school choice with randomly chosen tiebreak order, NBA draft.
- The unfairness we alluded to earlier has been perceived in real-world examples.

Order symmetry

- We define order symmetry, an average-case fairness concept in this ordinal setting: each agent has equal chance of getting their first choice, equal chance of their second item, etc.
- Formally, let P be a probability measure on preferences. We say a deterministic or fractional assignment mechanism is order symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is a concept of fairness at a point in time after the roles of the agents in the mechanism have been determined but before preferences are known.
- We call a randomized mechanism *ex post order symmetric* if it is described as a lottery over order symmetric deterministic mechanisms.

Order symmetry

- We define order symmetry, an average-case fairness concept in this ordinal setting: each agent has equal chance of getting their first choice, equal chance of their second item, etc.
- Formally, let P be a probability measure on preferences. We say a deterministic or fractional assignment mechanism is order symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is a concept of fairness at a point in time after the roles of the agents in the mechanism have been determined but before preferences are known.
- We call a randomized mechanism *ex post order symmetric* if it is described as a lottery over order symmetric deterministic mechanisms.

Order symmetry

- We define order symmetry, an average-case fairness concept in this ordinal setting: each agent has equal chance of getting their first choice, equal chance of their second item, etc.
- Formally, let P be a probability measure on preferences. We say a deterministic or fractional assignment mechanism is order symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is a concept of fairness at a point in time after the roles of the agents in the mechanism have been determined but before preferences are known.
- We call a randomized mechanism *ex post order symmetric* if it is described as a lottery over order symmetric deterministic mechanisms.

Order symmetry

- We define order symmetry, an average-case fairness concept in this ordinal setting: each agent has equal chance of getting their first choice, equal chance of their second item, etc.
- Formally, let P be a probability measure on preferences. We say a deterministic or fractional assignment mechanism is order symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is a concept of fairness at a point in time after the roles of the agents in the mechanism have been determined but before preferences are known.
- We call a randomized mechanism *ex post order symmetric* if it is described as a lottery over order symmetric deterministic mechanisms.

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Probability measures on preference profiles

- We need some restriction: if P has all weight on unanimous preferences, we can't satisfy order symmetry.
- P is anonymous if the identity of agents is irrelevant and neutral if the identity of objects is irrelevant.
- P is fully symmetric if both anonymous and neutral.
- Famous fully symmetric P:
  - IC (independent agents, each choosing uniform permutation of objects)
  - Uniform distribution on any class defined without singling out objects or agents, e.g. single peaked
  - IAC and other urn models
- Mallows preferences are NOT fully symmetric (not neutral).

Example (TTC fairer than SD under IC)

- Consider agents a_1, a_2, a_3 and objects o_1, o_2, o_3 under IC probability measure on preferences. The expected rank distribution matrix for SD with picking order a_1, a_2, a_3 is

$$\begin{bmatrix} 1 & 0 & 0 \\ 2/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}.$$

- For TTC with initial endowment $a_1 \leftarrow o_1, a_2 \leftarrow o_2, a_3 \leftarrow o_3$ we have

$$\begin{bmatrix} 2/3 & 2/9 & 1/9 \\ 2/3 & 2/9 & 1/9 \\ 2/3 & 2/9 & 1/9 \end{bmatrix}.$$

Order symmetry is a weakening of anonymity

No discrete mechanism can be anonymous. However:

Theorem

If A is an anonymous fractional assignment mechanism and P is an anonymous probability measure then A satisfies order symmetry with respect to P.

Huge difference between RSD and TTC

Theorem

Let P be a fully symmetric measure. Then TTC with any fixed endowment is order symmetric. Thus **TTC** *is ex post order symmetric.*

Theorem

Let P be a probability measure. The only way **RSD** *can be ex post order symmetric with respect to P is if P is supported on profiles in which all agents have different top choices.*

Example (TTC under cyclic permutation $1 \to 2 \to 3 \to 1$)

agent	profile	permute A	permute A and O
a_1	$o_1 \succ o_2 \succ o_3$	$o_2 \succ o_1 \succ o_3$	$o_3 \succ o_2 \succ o_1$
a_2	$o_1 \succ o_3 \succ o_2$	$o_1 \succ o_2 \succ o_3$	$o_2 \succ o_3 \succ o_1$
a_3	$o_2 \succ o_1 \succ o_3$	$o_1 \succ o_3 \succ o_2$	$o_2 \succ o_1 \succ o_3$
	endowment		
a_1	o_1	o_3	o_1
a_2	o_2	o_1	o_2
a_3	o_3	o_2	o_3
	allocation		
a_1	o_1	o_2	o_3
a_2	o_3	o_1	o_2
a_3	o_2	o_3	o_1
	rank		
a_1	1	1	1
a_2	2	1	1
a_3	1	2	2

Order symmetry is compatible with ordinal efficiency

- Ordinal efficiency is a strengthening of ex post Pareto efficiency that makes sense for fractional mechanisms.
- We defined a (computationally intractable) mechanism called **RMM-RA**.

Theorem

RMM-RA *is ex ante ordinally efficient and ex post order symmetric with respect to every fully symmetric* P.

Open question: Is the same true of Probabilistic Serial (Bogomolnaia and Moulin 2001)? In other words, can it be realized as a lottery over order symmetric discrete mechanisms, without losing its nice properties?

What if we don't have order symmetry?

- It may be desirable in some applications (e.g. sports draft) to avoid order symmetry.
- In any case it is useful to be able to quantify the deviation from order symmetry.
- For our basic computations we use the normalized gap in Borda welfare (linear utilities) between best-off and worst-off agent, in expectation over P.
- This is of course zero for order symmetric mechanisms.
- Other choices of how to measure unfairness are possible.

What if we don't have order symmetry?

- It may be desirable in some applications (e.g. sports draft) to avoid order symmetry.
- In any case it is useful to be able to quantify the deviation from order symmetry.
- For our basic computations we use the normalized gap in Borda welfare (linear utilities) between best-off and worst-off agent, in expectation over P.
- This is of course zero for order symmetric mechanisms.
- Other choices of how to measure unfairness are possible.

What if we don't have order symmetry?

- It may be desirable in some applications (e.g. sports draft) to avoid order symmetry.
- In any case it is useful to be able to quantify the deviation from order symmetry.
- For our basic computations we use the normalized gap in Borda welfare (linear utilities) between best-off and worst-off agent, in expectation over P.
- This is of course zero for order symmetric mechanisms.
- Other choices of how to measure unfairness are possible.

What if we don't have order symmetry?

- It may be desirable in some applications (e.g. sports draft) to avoid order symmetry.
- In any case it is useful to be able to quantify the deviation from order symmetry.
- For our basic computations we use the normalized gap in Borda welfare (linear utilities) between best-off and worst-off agent, in expectation over P.
- This is of course zero for order symmetric mechanisms.
- Other choices of how to measure unfairness are possible.

What if we don't have order symmetry?

- It may be desirable in some applications (e.g. sports draft) to avoid order symmetry.
- In any case it is useful to be able to quantify the deviation from order symmetry.
- For our basic computations we use the normalized gap in Borda welfare (linear utilities) between best-off and worst-off agent, in expectation over P.
- This is of course zero for order symmetric mechanisms.
- Other choices of how to measure unfairness are possible.

Order bias under sincere Mallows preferences for 4 mechanisms

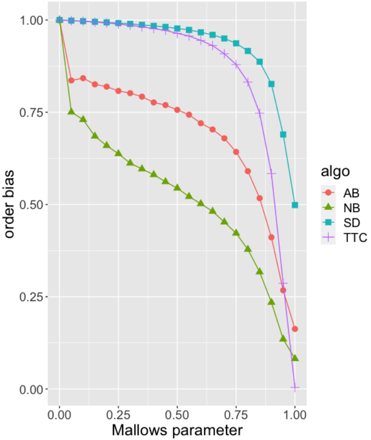

Figure: Mallows preferences, Borda order bias, $n = 64$, sample size 10000

Some questions

- Can Probabilistic Serial be decomposed as a lottery over order symmetric discrete mechanisms?
- Does TTC dominate SD with respect to order bias for every measure P?
- Is order symmetry compatible with other properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models? In particular how does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness idea useful more generally (e.g. fairness in AI/ML)?

Some questions

- Can Probabilistic Serial be decomposed as a lottery over order symmetric discrete mechanisms?
- Does TTC dominate SD with respect to order bias for every measure P?
- Is order symmetry compatible with other properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models? In particular how does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness idea useful more generally (e.g. fairness in AI/ML)?

Some questions

- Can Probabilistic Serial be decomposed as a lottery over order symmetric discrete mechanisms?
- Does TTC dominate SD with respect to order bias for every measure P?
- Is order symmetry compatible with other properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models? In particular how does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness idea useful more generally (e.g. fairness in AI/ML)?

Some questions

- Can Probabilistic Serial be decomposed as a lottery over order symmetric discrete mechanisms?
- Does TTC dominate SD with respect to order bias for every measure P?
- Is order symmetry compatible with other properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models? In particular how does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness idea useful more generally (e.g. fairness in AI/ML)?

Some questions

- Can Probabilistic Serial be decomposed as a lottery over order symmetric discrete mechanisms?
- Does TTC dominate SD with respect to order bias for every measure P?
- Is order symmetry compatible with other properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models? In particular how does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness idea useful more generally (e.g. fairness in AI/ML)?

Extensions

- After submission we were informed that in the special case of IC, the concept of order symmetry was already introduced in the unpublished PhD thesis of Xinghua Long (TAMU 2016); see Long & Velez (arXiv 2021).
- Previous audiences have alerted us to papers that are somewhat relevant: Harless & Manjunath, International Economic Review 2018; Pycia & Ünver, Theoretical Economics 2017.
- This is an idea whose time has come, and should be investigated in other social choice models.
- Geoff Pritchard and I are close to finishing a detailed analysis of the asymptotic distribution of the rank of the item obtained by each agent for Boston algorithms under IC.