Order Symmetry: a new fairness criterion for assignment mechanisms

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

AMMCS SPECIAL SESSION ON DECISIONS AND FAIRNESS, 2023-08-14

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 bijection (matching) between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.

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 bijection (matching) between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.

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 bijection (matching) between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.

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 bijection (matching) between A and O.
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.

Serial Dictatorship mechanism

- Serial Dictatorship (SD): 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.

Serial Dictatorship mechanism

- Serial Dictatorship (SD): 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.

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.
- Also strategyproof, Pareto efficient and easy to implement.

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.
- Also strategyproof, Pareto efficient and easy to implement.

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.
- Also strategyproof, Pareto efficient and easy to implement.

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.
- Also strategyproof, Pareto efficient and easy to implement.

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.
- Also strategyproof, Pareto efficient and easy to implement.

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.
- Also strategyproof, Pareto efficient and easy to implement.

Randomized versions

- We can choose the initial order randomly (obtaining the randomized mechanism **RSD**) or the initial endowment randomly (obtaining **TTC-RE**).
- Note that this is a (uniform) lottery over deterministic mechanisms.

Randomized versions

- We can choose the initial order randomly (obtaining the randomized mechanism **RSD**) or the initial endowment randomly (obtaining **TTC-RE**).
- Note that this is a (uniform) lottery over deterministic mechanisms.

SD and **TTC** are the same

- Abdulkadiroğlu and Sönmez (1998): **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante* and in expectation, and thus equally fair.
- Or are they?

SD and **TTC** are the same

- Abdulkadiroğlu and Sönmez (1998): **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante* and in expectation, and thus equally fair.
- Or are they?

SD and **TTC** are the same

- Abdulkadiroğlu and Sönmez (1998): **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two mechanisms are indistinguishable *ex ante* and in expectation, and thus equally fair.
- Or are they?

Giving away a whole apple or banana to Alice and Bob

Alice	Bob	SD	SD ranks	TTC	TTC ranks
$a \succ b$	$a \succ b$	$A : a, B : b$	$A : 1, B : 2$	$A : a, B : b$	$A : 1, B : 2$
$a \succ b$	$b \succ a$	$A : a, B : b$	$A : 1, B : 1$	$A : a, B : b$	$A : 1, B : 1$
$b \succ a$	$a \succ b$	$A : b, B : a$	$A : 1, B : 1$	$A : b, B : a$	$A : 1, B : 1$
$b \succ a$	$b \succ a$	$A : b, B : a$	$A : 1, B : 2$	$A : a, B : b$	$A : 2, B : 1$

Table: Alice, Bob and their fruits. We assume the initial endowment is $A : a, B : b$ for TTC, and Alice chooses first in SD.

Universal fairness

- Once the randomness is realized in **SD**, some agents may be advantaged by their very role in the mechanism, independent of preferences.
- This is not obviously so for **TTC**.
- We consider universal fairness guarantees that hold for all realizations.
- But there is no way to achieve universal fairness in the worst case.

Universal fairness

- Once the randomness is realized in **SD**, some agents may be advantaged by their very role in the mechanism, independent of preferences.
- This is not obviously so for **TTC**.
- We consider universal fairness guarantees that hold for all realizations.
- But there is no way to achieve universal fairness in the worst case.

Universal fairness

- Once the randomness is realized in **SD**, some agents may be advantaged by their very role in the mechanism, independent of preferences.
- This is not obviously so for **TTC**.
- We consider universal fairness guarantees that hold for all realizations.
- But there is no way to achieve universal fairness in the worst case.

Universal fairness

- Once the randomness is realized in **SD**, some agents may be advantaged by their very role in the mechanism, independent of preferences.
- This is not obviously so for **TTC**.
- We consider universal fairness guarantees that hold for all realizations.
- But there is no way to achieve universal fairness in the worst case.

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 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 weakening of anonymity, which itself can't be satisfied for deterministic house allocation anyway. It can also be thought of as fairness under uncertainty on preferences.
- We call a randomized mechanism *universally order symmetric* if it can be realized 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 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 weakening of anonymity, which itself can't be satisfied for deterministic house allocation anyway. It can also be thought of as fairness under uncertainty on preferences.
- We call a randomized mechanism *universally order symmetric* if it can be realized 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 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 weakening of anonymity, which itself can't be satisfied for deterministic house allocation anyway. It can also be thought of as fairness under uncertainty on preferences.
- We call a randomized mechanism *universally order symmetric* if it can be realized 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 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 weakening of anonymity, which itself can't be satisfied for deterministic house allocation anyway. It can also be thought of as fairness under uncertainty on preferences.
- We call a randomized mechanism *universally order symmetric* if it can be realized as a lottery over order symmetric deterministic mechanisms.

Huge difference between RSD and TTC

Theorem

Let P be an anonymous and neutral probability measure. Then **TTC** *is universally order symmetric with respect to P.*

Theorem

Let P be a probability measure. Then **RSD** *can be universally order symmetric with respect to P only if P is supported on the contention-free profiles, in which all agents have different top choices.*

The order symmetry lens

- Order symmetry can be achieved without tradeoffs in many cases.
- We are studying Boston mechanisms and have some partial results.
- Numerical simulation shows that simple tweaks (e.g. reverse the tiebreak order after the first round) can get us much closer to order symmetry.
- In the school choice realm, lack of order symmetry of some Boston mechanisms has been intuited by parents and officials.

The order symmetry lens

- Order symmetry can be achieved without tradeoffs in many cases.
- We are studying Boston mechanisms and have some partial results.
- Numerical simulation shows that simple tweaks (e.g. reverse the tiebreak order after the first round) can get us much closer to order symmetry.
- In the school choice realm, lack of order symmetry of some Boston mechanisms has been intuited by parents and officials.

The order symmetry lens

- Order symmetry can be achieved without tradeoffs in many cases.
- We are studying Boston mechanisms and have some partial results.
- Numerical simulation shows that simple tweaks (e.g. reverse the tiebreak order after the first round) can get us much closer to order symmetry.
- In the school choice realm, lack of order symmetry of some Boston mechanisms has been intuited by parents and officials.

The order symmetry lens

- Order symmetry can be achieved without tradeoffs in many cases.
- We are studying Boston mechanisms and have some partial results.
- Numerical simulation shows that simple tweaks (e.g. reverse the tiebreak order after the first round) can get us much closer to order symmetry.
- In the school choice realm, lack of order symmetry of some Boston mechanisms has been intuited by parents and officials.

Some open questions

- Can Probabilistic Serial be realized via a universally order-symmetric randomized mechanism?
- What can be said about order symmetry in other allocation models?
- What happens if we only require partial symmetry (e.g. between agents on the same side in 2-sided matching)?
- Is this average-case fairness idea useful more generally (e.g. group fairness in AI/ML)?

Some open questions

- Can Probabilistic Serial be realized via a universally order-symmetric randomized mechanism?
- What can be said about order symmetry in other allocation models?
- What happens if we only require partial symmetry (e.g. between agents on the same side in 2-sided matching)?
- Is this average-case fairness idea useful more generally (e.g. group fairness in AI/ML)?

Some open questions

- Can Probabilistic Serial be realized via a universally order-symmetric randomized mechanism?
- What can be said about order symmetry in other allocation models?
- What happens if we only require partial symmetry (e.g. between agents on the same side in 2-sided matching)?
- Is this average-case fairness idea useful more generally (e.g. group fairness in AI/ML)?

Some open questions

- Can Probabilistic Serial be realized via a universally order-symmetric randomized mechanism?
- What can be said about order symmetry in other allocation models?
- What happens if we only require partial symmetry (e.g. between agents on the same side in 2-sided matching)?
- Is this average-case fairness idea useful more generally (e.g. group fairness in AI/ML)?

Stuff there wasn't time for today

- Quantifying lack of order symmetry, and comparing mechanisms
- Substantial numerical simulation results
- Order symmetry is compatible with ex ante ordinal efficiency
- Check out the original preprint (currently under heavy revision): https://osf.io/preprints/socarxiv/xt37c/

Stuff there wasn't time for today

- Quantifying lack of order symmetry, and comparing mechanisms
- Substantial numerical simulation results
- Order symmetry is compatible with ex ante ordinal efficiency
- Check out the original preprint (currently under heavy revision): https://osf.io/preprints/socarxiv/xt37c/

Stuff there wasn't time for today

- Quantifying lack of order symmetry, and comparing mechanisms
- Substantial numerical simulation results
- Order symmetry is compatible with ex ante ordinal efficiency
- Check out the original preprint (currently under heavy revision): https://osf.io/preprints/socarxiv/xt37c/

Stuff there wasn't time for today

- Quantifying lack of order symmetry, and comparing mechanisms
- Substantial numerical simulation results
- Order symmetry is compatible with ex ante ordinal efficiency
- Check out the original preprint (currently under heavy revision): https://osf.io/preprints/socarxiv/xt37c/