Order Symmetry: a new fairness criterion for assignment rules

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

MATCH-UP, Vienna, 2022-08-25

Example (Birthday party drama)

- Each child gets exactly one toy.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Each child gets exactly one toy.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Each child gets exactly one toy.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Example (Birthday party drama)

- Each child gets exactly one toy.
- However, you have no idea as to the preferences of the children.
- How should you allocate the toys?

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Formal model

- A: finite set of n agents; O: set of n objects.
- Set of all strict linear orders of objects: $L(O)$.
- Set of all profiles is $X := L(O)^A$.
- A matching is a bijection $A \to O$; the set of all such is $M(A, O)$.
- The house allocation problem: find a matching rule $f : X \to M(A, O)$.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Informal description

- Given strict ordinal preferences of agents over objects, match each agent with an object!
- Common applications: students to dorm rooms, military to overseas postings, professors to offices.
- Closely related problems: school choice, multi-unit assignment.
- Key standard axiomatic properties:
  - Pareto efficiency: can't help someone without hurting someone else;
  - Anonymity: identities of agents are irrelevant to the matching;
  - Strategyproofness: no agent ever has incentive to lie about their preferences.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

This rule is strategyproof, Pareto efficient and easy to implement.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

This rule is strategyproof, Pareto efficient and easy to implement.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

This rule is strategyproof, Pareto efficient and easy to implement.

Commonly used solution: serial dictatorship (SD)

- Fix an exogenous order on agents.
- Let them choose in turn, according to this order, their favorite remaining object.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Another solution: Top Trading Cycles (TTC)

- Each agent is given an exogenous initial allocation (the endowment).
- Each agent points to the owner of their favorite object.
- This creates a directed graph which must have at least one cycle.
- We can and do resolve all cycles by giving each agent in each cycle their desired object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is strategyproof, Pareto efficient and easy to implement.

Which seems fairer?

- Under SD, 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.
- The same bias is not obviously true of TTC.

Which seems fairer?

- Under SD, 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.
- The same bias is not obviously true of TTC.

Which seems fairer?

- Under SD, 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.
- The same bias is not obviously true of TTC.

Order symmetry

- We define *order symmetry*: each agent has equal chance of getting their first choice, equal chance of their second, etc.
- Formally, let P be a probability measure on preferences. We say a rule is order-symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is an average-case fairness concept, not a worst-case axiomatic property.
- It considers procedural fairness rather than an outcome fairness.
- It is a weakening of anonymity (which can never be satisfied by deterministic matchings of our type).

Order symmetry

- We define *order symmetry*: each agent has equal chance of getting their first choice, equal chance of their second, etc.
- Formally, let P be a probability measure on preferences. We say a rule is order-symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is an average-case fairness concept, not a worst-case axiomatic property.
- It considers procedural fairness rather than an outcome fairness.
- It is a weakening of anonymity (which can never be satisfied by deterministic matchings of our type).

Order symmetry

- We define *order symmetry*: each agent has equal chance of getting their first choice, equal chance of their second, etc.
- Formally, let P be a probability measure on preferences. We say a rule is order-symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is an average-case fairness concept, not a worst-case axiomatic property.
- It considers procedural fairness rather than an outcome fairness.
- It is a weakening of anonymity (which can never be satisfied by deterministic matchings of our type).

Order symmetry

- We define *order symmetry*: each agent has equal chance of getting their first choice, equal chance of their second, etc.
- Formally, let P be a probability measure on preferences. We say a rule is order-symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is an average-case fairness concept, not a worst-case axiomatic property.
- It considers procedural fairness rather than an outcome fairness.
- It is a weakening of anonymity (which can never be satisfied by deterministic matchings of our type).

Order symmetry

- We define *order symmetry*: each agent has equal chance of getting their first choice, equal chance of their second, etc.
- Formally, let P be a probability measure on preferences. We say a rule is order-symmetric with respect to P if the expected rank distribution matrix with respect to P has all rows equal.
- This is an average-case fairness concept, not a worst-case axiomatic property.
- It considers procedural fairness rather than an outcome fairness.
- It is a weakening of anonymity (which can never be satisfied by deterministic matchings of our type).

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

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

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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, 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 it is 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 (anonymous, not neutral).

Huge difference between SD and TTC

Theorem

Let P be a fully symmetric probability measure. Then TTC with any fixed endowment is order-symmetric.

Corollary

Order symmetry is compatible with strategyproofness and Pareto efficiency.

Theorem

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

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 rules.
- 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 rules.
- 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 rules.
- 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 rules.
- 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 rules.
- Other choices of how to measure unfairness are possible.

Order bias under Mallows preferences

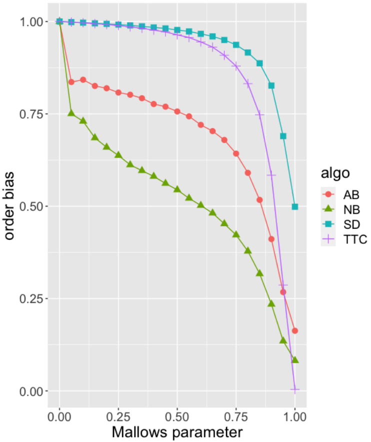

Figure: Mallows preferences, Borda order bias, $n = 64$, sample size 10000

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Another solution: (Naive) Boston

- Choose an exogenous tiebreaking order.
- In round i, all remaining agents bid for their ith preference.
- Use the tiebreaking order to decide who gets an object.
- Continue with the remaining agents, after removing the satisfied agents and their objects.

This rule is not strategyproof, but is Pareto efficient and easy to implement, and usually has better welfare.

Reversing version

- Naive Boston has bias toward those earlier in the tiebreaking order.
- To reduce this, we might reverse the order at rounds 2 and higher.
- There is no change in social welfare if we do this.
- Without the conceptual idea of order bias, we would probably not notice this variant.

Reversing version

- Naive Boston has bias toward those earlier in the tiebreaking order.
- To reduce this, we might reverse the order at rounds 2 and higher.
- There is no change in social welfare if we do this.
- Without the conceptual idea of order bias, we would probably not notice this variant.

Reversing version

- Naive Boston has bias toward those earlier in the tiebreaking order.
- To reduce this, we might reverse the order at rounds 2 and higher.
- There is no change in social welfare if we do this.
- Without the conceptual idea of order bias, we would probably not notice this variant.

Reversing version

- Naive Boston has bias toward those earlier in the tiebreaking order.
- To reduce this, we might reverse the order at rounds 2 and higher.
- There is no change in social welfare if we do this.
- Without the conceptual idea of order bias, we would probably not notice this variant.

Order bias under Mallows (0.6) preferences

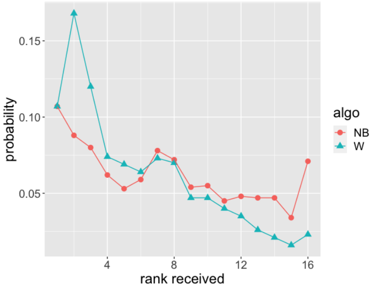

Figure: Boston versus its reversing version, $n = 16$, fate of last agent

Clarifying a famous result

- Let **RSD** be the randomized matching rule that chooses a random order on agents and then runs SD with that order.
- Let **TTC** be the randomized matching rule that chooses a random endowment and then runs TTC with that endowment.
- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two rules are indistinguishable *ex ante*, and give the same expected (fractional) assignment.
- However we have seen that one is explicitly implemented as a lottery over order-symmetric matching rules, and the other is not: *ex post* fairness behavior is very different.

Clarifying a famous result

- Let **RSD** be the randomized matching rule that chooses a random order on agents and then runs SD with that order.
- Let **TTC** be the randomized matching rule that chooses a random endowment and then runs TTC with that endowment.
- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two rules are indistinguishable *ex ante*, and give the same expected (fractional) assignment.
- However we have seen that one is explicitly implemented as a lottery over order-symmetric matching rules, and the other is not: *ex post* fairness behavior is very different.

Clarifying a famous result

- Let **RSD** be the randomized matching rule that chooses a random order on agents and then runs SD with that order.
- Let **TTC** be the randomized matching rule that chooses a random endowment and then runs TTC with that endowment.
- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two rules are indistinguishable *ex ante*, and give the same expected (fractional) assignment.
- However we have seen that one is explicitly implemented as a lottery over order-symmetric matching rules, and the other is not: *ex post* fairness behavior is very different.

Clarifying a famous result

- Let **RSD** be the randomized matching rule that chooses a random order on agents and then runs SD with that order.
- Let **TTC** be the randomized matching rule that chooses a random endowment and then runs TTC with that endowment.
- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two rules are indistinguishable *ex ante*, and give the same expected (fractional) assignment.
- However we have seen that one is explicitly implemented as a lottery over order-symmetric matching rules, and the other is not: *ex post* fairness behavior is very different.

Clarifying a famous result

- Let **RSD** be the randomized matching rule that chooses a random order on agents and then runs SD with that order.
- Let **TTC** be the randomized matching rule that chooses a random endowment and then runs TTC with that endowment.
- Abdulkadiroğlu and Sönmez (1998) proved that **RSD** and **TTC** give the same mapping from preference profiles to lotteries over assignments.
- The two rules are indistinguishable *ex ante*, and give the same expected (fractional) assignment.
- However we have seen that one is explicitly implemented as a lottery over order-symmetric matching rules, and the other is not: *ex post* fairness behavior is very different.

Some questions

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

Some questions

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

Some questions

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

Some questions

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

Some questions

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

Some questions

- Does TTC dominate SD with respect to order bias for every measure P?
- Can Probabilistic Serial be implemented as a lottery over order-symmetric matching rules?
- Is order symmetry compatible with other axiomatic properties, such as obvious strategyproofness?
- What can be said about order symmetry in other allocation models?
- How does the idea of order symmetry relate to other fairness criteria such as envy-freeness and EF1?
- Is this average-case fairness approach 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 (called "balancedness") was already introduced in the unpublished PhD thesis of Xinghua Long (TAMU 2016); see Long & Velez (arXiv 2021).
- Geoff Pritchard and I have (`https://arxiv.org/pdf/2205.15418.pdf`; under review) a detailed analysis of the asymptotic distribution under IC of the rank distribution matrix for several Boston algorithms.
- We believe this is an idea whose time has come, and should be investigated in other social choice models.

Order symmetry is a weakening of anonymity

No deterministic matching rule can be anonymous. However:

Theorem

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