

Order Symmetry: A New Fairness Criterion for Assignment Mechanisms*

Rupert Freeman, Geoffrey Pritchard, Mark Wilson

Abstract

We introduce a new fairness criterion, *order symmetry*, for assignment mechanisms that match n objects to n agents with ordinal preferences over the objects. An assignment mechanism is order symmetric with respect to some probability measure over preference profiles if every agent is equally likely to receive their favorite object, every agent is equally likely to receive their second favorite, and so on. When associated with a sufficiently symmetric probability measure, order symmetry is a relaxation of anonymity that, crucially, can be satisfied by discrete assignment mechanisms. Furthermore, it can be achieved without sacrificing other desirable axiomatic properties satisfied by existing mechanisms. In particular, we show that it can be achieved in conjunction with strategyproofness and ex post efficiency via the Top Trading Cycles mechanism (but not Serial Dictatorship). We additionally design a novel mechanism that is both order symmetric and ordinally efficient. The practical utility of order symmetry is substantiated by simulations on Impartial Culture and Mallows-distributed preferences for four common assignment mechanisms.

1 Introduction

Imagine you are hosting a birthday party for a group of children. You have bought 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 toys 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?

This toy (pun intended) problem (inspired by the title of (Thomson 2007)) is an instance of the *house allocation problem* (also known as the assignment problem or the one-sided matching problem). In that problem, n indivisible objects must be matched to n agents, each with an ordinal preference over the objects, and monetary transfers are not allowed. It models a number of real-world resource allocation settings such as assigning rooms to college students and assigning schools to students in a school district, and has been studied extensively in economics (Ergin 2000, Bogomolnaia and Moulin 2001, Sönmez and Ünver 2010), operations research (Saban and Sethuraman 2015, Crès and Moulin 2001, Bade 2020), and computer science (Abraham et al. 2004, Saban and Sethuraman 2013, Krysta et al. 2019, Gan et al. 2019). The model was introduced by Hylland and Zeckhauser (1979), who also mention assigning legislators to committees.

In the discrete setting as described above, some amount of unfairness is inherent in the problem. Consider the case where all agents have the same preference over the objects; one agent necessarily receives the best object, while another must receive the worst. For many traditional notions of fairness, this is the end of the story. For example, this instance already witnesses a violation of anonymity (requiring that agents are treated symmetrically), equal-treatment-of-equals (requiring that any agents with the same preference receive the same allocation), and envy-freeness (no agent should prefer another agent's allocation to her own) (Foley 1967).

Owing to this difficulty, much of the attention in both theory and practice has focused on randomized assignment mechanisms. Following the seminal work of Hylland and Zeckhauser (1979), typically the goal is to provide both ex ante and ex post guarantees, with the former applying in expectation over the mechanism's randomness, and the latter applying after the randomness has been realized (Bogomolnaia and Moulin 2001, Chen and Sönmez 2002, Nesterov 2017). Existing fairness properties are necessarily violated ex post, so are usually only considered in the ex ante sense, with ex post guarantees reserved for efficiency.

*Author ordering is alphabetical.

Perhaps the most common mechanism used in practice is *random serial dictatorship* (RSD), which randomly chooses an ordering of the agents and allows each agent to choose, in that order, her favorite remaining object. RSD is simple to execute and enjoys compelling theoretical justification: it is known to be strategyproof, ex ante anonymous, and ex post efficient (however, it is not ex ante efficient—indeed this combination of properties is ruled out by Zhou (1990)). It also produces the same lottery over assignments as the *Top Trading Cycles* (TTC) procedure, which endows each agent an object at random, and performs specified mutually beneficial trades of objects among agents (Abdulkadiroğlu and Sönmez 1998). TTC is known to produce the unique core outcome in the housing market model where each agent enters the mechanism already in possession of an object (Shapley and Scarf 1974, Roth and Postlewaite 1977). Bade (2020) has generalized the result of Abdulkadiroğlu and Sönmez (1998) to show that all strategyproof, Pareto optimal, and nonbossy mechanisms, when suitably symmetrized, coincide with RSD.

By the result of Abdulkadiroğlu and Sönmez (1998), RSD and TTC with random endowment give the same mapping from preference profiles to lotteries over assignments. After the agents report their preferences, the mechanism can randomly choose an ordering of the agents and simulate the iterative choosing (or, equivalently, randomly choose an initial endowment and run TTC). However, we argue that in many settings this framework is still not sufficient, since it is often desirable for the mechanism’s randomness to be realized *before* the agents reveal their preferences (Pycia and Ünver (2015) make a similar argument). Indeed, many real-world mechanisms operate in this way. For example, in the NBA draft, the exact order in which teams pick players is determined by a combination of random draw and team strength, but the ordering is decided upon before the draft occurs. When college housing is allocated via RSD, students will typically receive their position in the ordering before any preferences are revealed. This method has several advantages: it has lower communication complexity (each agent only needs to select one object in turn, not rank all of them), it allows agents to more easily reason about their decisions (Li 2017), and it de-emphasizes the randomization. For the latter point, note that a malicious actor who is able to influence the randomization process will wield less power if the randomization occurs before preferences are elicited.

For these reasons, we study randomized assignment mechanisms that select a discrete assignment mechanism at random *before* eliciting agent preferences. For RSD, this means randomly choosing an agent ordering, while for TTC a random endowment, or initial assignment, is chosen. When defined in this manner, the two mechanisms are no longer equivalent. It is clear that RSD remains highly unfair, in the sense that the agent first in line has a substantial advantage over the last. However, TTC admits no such obvious systematic bias toward any specific agent. While some initial endowments may be better than others, it is not a priori clear which agents will benefit, without observing the preferences.

To illustrate this point, let us return to the example of the children’s party. Running RSD amounts to randomly choosing an ordering of the children and having each child choose a toy in turn. This is fair in expectation because each child has equal probability of occupying each position in line. However, the realization of the mechanism’s randomness has a large influence over how happy each child will be with their assigned toy—the first child receives their top choice for sure, while the last is left to take whatever toy is left at the end. Under TTC, the mechanism’s randomness manifests in a preliminary random assignment of a toy to each child. Without knowing the preferences of the children, it is impossible to say whether a given initial allocation is good for a particular child or not, or whether there is systematic bias toward one child. This hints at the possibility of an elusive, nontrivial ex post fairness notion for house allocation mechanisms (that is, a fairness notion that retains meaning when applied to deterministic mechanisms).

Our proposal is to consider average case fairness: for some probability measure on profiles, a mechanism is fair (with respect to that measure) if all agents do equally well *in expectation*. This nicely captures the intuition for why serial dictatorship and other mechanisms that rely heavily on a fixed choosing order or tiebreak order are unfair, since some agents do better than (or at least as well as) others systematically across all profiles. Since agents only provide ordinal preferences and not cardinal utilities, we measure how highly they rank the object they receive. Each agent should have equal probability of receiving their top object, their second object, and so on. When these probabilities are exactly equal for all agents according to some measure P on profiles, we say that the mechanism is *order symmetric* with respect to P , reflecting the intuition that agents are not systematically ordered by the mechanism. If probabilities are not exactly equal, then we can examine the discrepancy to yield a measure of unfairness that we call the *order bias* of the mechanism.

In independent and simultaneous work, Long and Velez (2021) (based on the results of Long (2016)) also consider average case fairness for house allocation and obtain complementary results to ours. They

define balancedness, which is equivalent to order symmetry with respect to the uniform measure on profiles, and show that TTC is (with an exception for three agents) the only mechanism that is balanced, efficient, and group strategyproof. However, they do not consider non-uniform measures on profiles, or ex ante properties of randomized mechanisms. Some other work also considers average case fairness. Pycia and Ünver (2014) show that their *brokered trading cycles* mechanisms can be more equitable (on average over profiles) to agents than TTC in an asymmetric house allocation problem in which the mechanism designer has preferences over allocations (for example, the mechanism designer prefers that a certain agent not be allocated a certain object); we do not consider such preferences in our work. Harless and Manjunath (2018) consider a model where agents have the same expected utility function over items but each is able to discover their exact utility for only one object. While the model is quite different from ours the conclusion is similar: TTC distributes expected utility (where the expectation is taken over the uncertainty in the utility functions) more evenly than serial dictatorship. Other work examining average case analysis in the one-sided matching setting (Filos-Ratsikas et al. 2014, Deng et al. 2017, Gao and Zhang 2019) has focused on social welfare approximation rather than achieving fairness. The work of Manshadi et al. (2021) is conceptually similar, in that they design a mechanism with guarantees on the minimum expected (ex ante) and the expected minimum (ex post) utility, where the expectation is taken with respect to randomness in the agents’ demands. However, the exact problem is quite different to ours, with agents arriving dynamically and demanding some quantity of a divisible good.

1.1 Why do we care about order symmetry?

Practitioners who are fully content with ex ante fairness guarantees may not be interested in the concept of order symmetry, since it deals with an ex post fairness criterion. However, we believe that there are many situations where order symmetry ought to be seriously considered. We now consider several issues, in no particular order of importance.

First, a randomized mechanism that is fair ex ante but unfair ex post assigns undue emphasis to the randomization component. Unlucky agents can be condemned to receiving bad assignments in a large fraction of all possible preference profiles. Our children’s party example may be contrived, but similar concerns appear in practice. Pathak and Sethuraman (2011) (also see (Long and Velez 2021)) describe how randomly assigning New York City public school students a global priority number was perceived by participating parents to be less fair than using a random priority ordering for each school independently, even though the two mechanisms are ex ante equivalent. Even worse, an agent who can manipulate the source of randomness can realize large utility gains. Even if the randomization mechanism is uncompromised, its integrity may be called into question if, say, a powerful actor is drawn in a favorable position. Striving for ex post fairness reduces the importance of the randomness and mitigates these concerns.

Second, in many situations there may be resistance from users (e.g., in school choice situations) toward randomized mechanisms being used, for reasons of trust, replicability, algorithmic simplicity, transparency, and ex-post fairness. The same issue occurs in voting, where tiebreaking issues are important. When randomization is disallowed, ex post guarantees are the only ones available. It is consistent with the ideal of fairness that, in the absence of information about preferences, no agent should have a systematic advantage over another. All else being equal, a mechanism that satisfies order symmetry should be preferred to one that does not.

Third, it is important to understand and quantify the deviation from order symmetry in settings where we cannot, or do not want to, achieve full order symmetry. For example, the picking order for some sports team drafts is partly or completely determined by the reverse of the success of the teams in the previous year, so there is an attempt made to equalize teams by using inherent bias toward the first teams in the picking order.

1.2 Our contribution, and outline of the paper

We make three main theoretical contributions.

- (i) In Section 3 we define the concept of *order symmetry*, a natural and novel fairness guarantee for assignment mechanisms. With respect to a sufficiently symmetric measure on profiles, order symmetry is a weakening of anonymity which, unlike anonymity, can be satisfied by discrete assignment mechanisms.

- (ii) In Section 4.2, we show that the well-known Top Trading Cycles (TTC) mechanism is order symmetric with respect to any fully symmetric probability measure on profiles (that is, one that is symmetric with respect to both agents and objects). The result separates TTC from RSD, suggesting an advantage of the former over the latter, despite both inducing the same fractional assignment.
- (iii) In Section 4.3, we show, by defining the **Rank Maximal Matching** mechanism, that ex ante efficiency is compatible with order symmetry when equipped with a fully symmetric probability measure.

Section 2 carefully covers basic background definitions and notations, and we conclude with possible directions for future research in Section 5. In addition to the material presented in the main text, in Appendix A.1 we define measures of order bias, or failure of order symmetry, and in Appendix A.2 we explore order bias of four standard mechanisms with respect to Impartial Culture and Mallows-distributed preferences. Our simulation results confirm that TTC is a better choice than RSD from the perspective of order bias. In the Mallows setting, we find that the Naive Boston mechanism incurs strikingly low order bias, even relative to TTC. Despite lacking theoretical justification, our results suggest that it may be a strong choice for reducing unfairness stemming from the mechanism’s randomness in settings where some objects are inherently more valuable than others.

2 Preliminaries

We now collect some basic terminology.

Definition 2.1. Let n be a positive integer and let $O = \{o_1, \dots, o_n\}$ be an ordered set of **objects** having cardinality n and $A = \{a_1, \dots, a_n\}$ an ordered set of **agents**. Each agent a_k has a strict linear order $\succ_k$ over the objects, called its **preference order**. This yields a **preference profile** which we write $\pi = (\succ_1, \dots, \succ_n) = (\pi(1), \dots, \pi(n))$. We write the preference order of agent k as $(\pi(k)_1, \dots, \pi(k)_n)$. We let $\Pi(A, O)$ denote the set of all preference profiles.

Definition 2.2 (assignments). A **fractional assignment** α is a mapping taking each agent to a probability measure on O such that each object is assigned to exactly one agent in expectation. Note that each assignment can be represented by a doubly stochastic matrix with rows indexed by agents and columns by objects. A fractional assignment is **discrete** if every entry of the matrix is either 0 or 1. We use “assignment” to mean a fractional assignment unless otherwise stated.

Definition 2.3 (assignment mechanisms). A **fractional assignment mechanism** $\mathcal{A}$ is a function associating a fractional assignment $\mathcal{A}(\pi)$ with each profile π . A fractional assignment mechanism is called **discrete** if it always outputs a discrete assignment. When we say “assignment mechanism” we will mean a fractional assignment mechanism unless otherwise clear from context.

A **randomized assignment mechanism** $\mathbf{A}$ is a lottery (probability measure) over discrete assignment mechanisms (we denote randomized mechanisms by bold letters for clarity). The set of discrete mechanisms occurring with positive probability is called the **support** of $\mathbf{A}$. We write q_j for the probability of discrete algorithm $\mathcal{A}_j$.

Each randomized assignment mechanism $\mathbf{A}$ induces a fractional assignment mechanism $\mathcal{A}$, via $\mathcal{A}(\pi) = \sum_j q_j \mathcal{A}_j(\pi)$. We say that $\mathbf{A}$ **implements** $\mathcal{A}$, or that $\mathcal{A}$ has **implementation** $\mathbf{A}$, and call the expression $\sum_j q_j \mathcal{A}_j$ a **decomposition** of $\mathcal{A}$.

Note that the induced fractional mechanism is simply the expectation with respect to the lottery of the randomized mechanism.

By the Birkhoff-von Neumann theorem (Birkhoff 1946, von Neumann 1953), every fractional assignment can be expressed as a convex combination of discrete assignments. In Appendix G we describe how any fractional assignment mechanism can be implemented by a randomized assignment mechanism. Note that a given assignment mechanism may have more than one decomposition.

Example 2.4. The **Serial Dictatorship** (SD) mechanism works as follows: fix an exogenous order ρ on the agents, and let them choose according to ρ their highest ranked object from those remaining. The **Random Serial Dictatorship** (RSD) mechanism works by first selecting the order ρ uniformly at random, and then running SD with that order. Thus, **RSD** is a randomized assignment mechanism, with support indexed by the set of all $n!$ orders of agents and the lottery being uniform. The induced

fractional assignment mechanism *RSD* is given by

$$RSD(\pi) = \frac{1}{n!} \sum_{\rho} SD_{\rho}(\pi)$$

where SD_{ρ} is *SD* with exogenous order ρ .

Let us also define the **Top Trading Cycles** (TTC) algorithm, originally presented by Shapley and Scarf (1974) and attributed to David Gale. Fix an initial assignment of objects to agents; this initial assignment is often known as an endowment. Begin by having each agent i point to the agent who is currently assigned to agent i 's most preferred object. It is clear that there must be at least one cycle in this graph (potentially of length one, if some agent i is already assigned their first choice). Choose one cycle and implement the trade indicated by that cycle by reallocating each object to the agent pointing to it. Then, remove all agents and objects involved in the trade. If any agents remain, have agents once again point to their most-preferred remaining object and continue clearing cycles. The **Top Trading Cycles with Random Endowment** (TTC-RE) mechanism is the randomized allocation mechanism that first selects an initial assignment and then runs TTC with that initial assignment.

It is known that the induced fractional assignment, TTC-RE, is equal to RSD. That is, $TTC-RE(\pi) = RSD(\pi)$ for all profiles π . In fact, an even stronger equivalence holds. For every profile π , **RSD** and **TTC-RE** induce the same lottery over assignments in the sense that choosing an ordering at random and running SD with that ordering is exactly equivalent to choosing an initial endowment at random and running TTC with that endowment (Abdulkadiroğlu and Sönmez 1998, Bade 2020).

We now introduce a number of properties that a fractional assignment mechanism may satisfy. For a property $\langle P \rangle$, we say that a randomized allocation mechanism $\mathcal{A}$ satisfies $\langle P \rangle$ ex ante if the fractional allocation $\mathcal{A}$ that it implements satisfies $\langle P \rangle$. We say that $\mathcal{A}$ satisfies $\langle P \rangle$ ex post if every discrete assignment mechanism in the support of $\mathcal{A}$ satisfies $\langle P \rangle$. Note that we are using the terms ex ante and ex post with respect to the randomness in the mechanism, and that this randomness occurs *before* preferences are realized.

It is intuitively clear that RSD treats agents symmetrically in a way that SD does not. This intuition is captured by notion of anonymity, which requires that the output of a mechanism should depend only on the preference orders and not on agent identities.

We define anonymity in terms of group actions. Let G denote the group $\text{Sym}(A)$ of all permutations of the set of agents. Since we have ordered the set $A = \{a_1, \dots, a_n\}$ we may identify G as usual with the symmetric group S_n . The group G acts on $\Pi(A, O)$ via $(g \cdot \pi)(a) = \pi(g^{-1}(a))$, and on assignments by $(g \cdot \alpha)(a) = \alpha(g^{-1}(a))$. We provide additional mathematical background on group actions in Appendix H for the interested reader.

Definition 2.5 (anonymity). *Assignment mechanism $\mathcal{A}$ is **anonymous** if and only if for each $g \in G$ and each $\pi \in \Pi(A, O)$, $\mathcal{A}(g \cdot \pi) = g \cdot \mathcal{A}(\pi)$.*

An example illustrating the requirement is provided in Appendix I. Non-anonymous mechanisms privilege some agents over others, and anonymity is a widely used fairness criterion. However, anonymity can only be achieved ex ante.

Proposition 2.6. *Discrete assignment mechanisms are never anonymous. Equivalently, randomized assignment mechanisms are never ex-post anonymous.*

Proof. Let $g \in G$ not be the identity, and let π be a profile where all agents have the same preference. Then $g \cdot \pi = \pi$ and so $\mathcal{A}(g \cdot \pi) = \mathcal{A}(\pi) \neq g \cdot \mathcal{A}(\pi)$. $\square$

Given a discrete assignment mechanism, it is easy to construct an ex ante anonymous randomized assignment mechanism that randomizes the agents' roles (Bade 2020). We provide details of the construction in Appendix H.2.

We will be interested in the following notion of economic efficiency. An agent ordinally prefers assignment α to assignment α' if, for every $k \in \{1, \dots, n\}$, she is allocated a greater share of her k most preferred objects in α than α' . Ordinal efficiency requires that a mechanism never outputs an assignment α when all agents weakly ordinally prefer a different assignment α' and at least one agent strictly ordinally prefers α' to α . Ordinal efficiency is often referred to as stochastic dominance efficiency (SD-efficiency) in the literature.

Definition 2.7 (efficiency). Let α, α' be assignments. We write $\alpha \succeq_i^{ord} \alpha'$, that is, agent i ordinally prefers assignment α to α' , if $\sum_{o_j \in \{o_k : o_k \succ_o\}} \alpha_{i,j} \geq \sum_{o_j \in \{o_k : o_k \succ_o\}} \alpha'_{i,j}$ for all $o \in O$. We write $\alpha \succ_i^{ord} \alpha'$ if $\alpha \succeq_i^{ord} \alpha'$ and not $\alpha' \succeq_i^{ord} \alpha$.

We say that α is **ordinally efficient** if there does not exist an α' with $\alpha' \succeq_i^{ord} \alpha$ for all agents a_i , and $\alpha' \succ_k^{ord} \alpha$ for some agent a_k . We say that an assignment mechanism $\mathcal{A}$ is **ordinally efficient** if $\mathcal{A}(\pi)$ is ordinally efficient for all $\pi \in \Pi(A, O)$.

Note that if we restrict α, α' to be discrete assignments, there is only one natural way to extend agents' preferences over objects to preferences over assignments. Thus, when discussing efficiency in a discrete sense, we will use the standard terminology **Pareto efficiency**, reserving ordinal efficiency for the fractional case.

It is known that ex ante ordinal efficiency implies ex post Pareto efficiency but the converse does not hold (Bogomolnaia and Moulin 2001).

Finally, we would like assignment mechanisms for which an agent can never improve her assignment by misreporting her preference to the mechanism. For a profile π , let π_{-i} denote the reports of all agents other than a_i .

Definition 2.8 (strategyproofness). Assignment mechanism $\mathcal{A}$ is **strategyproof** if, for all profiles $\pi \in \Pi(A, O)$, there does not exist a report $\succ'_i$ such that

$$\mathcal{A}(\succ'_i, \pi_{-i}) \succ_i^{ord} \mathcal{A}(\pi).$$

Note that the ordinal preference relation is complete when restricted to discrete assignments, so that if $\mathcal{A}$ is discrete then strategyproofness requires that $\mathcal{A}(\pi) \succeq_i^{ord} \mathcal{A}(\succ'_i, \pi_{-i})$ for all $\succ'_i$.

Theorem 2.9. If $\mathcal{A}$ is ex post strategyproof then it is ex ante strategyproof.

Proof. Fix a profile π and an agent a_i . For all j and all $\succ'_i$, it follows from ex post strategyproofness of $\mathcal{A}$ that for each discrete assignment mechanism $\mathcal{A}_j$ in the decomposition of $\mathcal{A}$, $\mathcal{A}_j(\pi) \succeq_i^{ord} \mathcal{A}_j(\succ'_i, \pi_{-i})$. Therefore, $\sum_{j=1}^m q_j \mathcal{A}_j(\pi) \succeq_i^{ord} \mathcal{A}_j(\succ'_i, \pi_{-i})$. $\square$

With respect to the properties discussed so far, **RSD** and **TTC-RE** are indistinguishable. Both are ex ante anonymous, ex post strategyproof, and ex post Pareto efficient (Shapley and Scarf 1974, Roth 1982). In the following section we introduce a property that distinguishes between the two.

3 Order symmetry

While anonymity is a reasonable requirement for assignment mechanisms that output fractional assignments, it is too strong for discrete assignment mechanisms, as Proposition 2.6 shows. In this section we introduce a new fairness criterion that is implied by anonymity and allows us to discriminate between discrete assignment mechanisms. We have motivated this concept above, and it is now time to define it formally.

3.1 Order symmetry definition

Definition 3.1. Let $\mathcal{A}$ be an assignment mechanism and π a profile. The **rank distribution** under mechanism $\mathcal{A}$ at profile π is the mapping $D_{\pi, \mathcal{A}}$ on $\{1, \dots, n\} \times \{1, \dots, n\}$ whose value at (r, j) is the probability that $\mathcal{A}$ assigns agent r the object it ranks as j th best.

Example 3.2. Consider the profile where agents a_1, a_2, a_3 have respective preferences over objects o_1, o_2, o_3 as follows.

$$\begin{aligned} a_1 : o_1 &\succ o_2 \succ o_3 \\ a_2 : o_1 &\succ o_3 \succ o_2 \\ a_3 : o_2 &\succ o_1 \succ o_3 \end{aligned}$$

The matrix representing the assignment made by SD with picking order a_1, a_2, a_3 is $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ where rows represent agents a_1, a_2, a_3 and columns represent objects o_1, o_2, o_3 . The matrix representing the assignment made by RSD is computed by averaging over all picking orders: $\begin{bmatrix} \frac{3}{6} & \frac{1}{6} & \frac{2}{6} \\ \frac{2}{6} & 0 & \frac{3}{6} \\ 0 & \frac{5}{6} & \frac{1}{6} \end{bmatrix}$. Finally consider

TTC with initial endowment that assigns object o_i to agent a_i for each $i \in \{1, 2, 3\}$. We verify that it gives the same assignment as SD. The top choice of agents a_1 and a_2 is object o_1 , so both agents point to a_1 . Agent a_3 has object o_2 as her top choice, so she points to agent a_2 . The only cycle in this graph is the cycle of length one in which a_1 points to herself, so we allocate o_1 to a_1 and remove them. Of the remaining objects, agent a_2 prefers object o_3 (so points to a_3) and agent a_3 prefers object o_2 (so points to a_2). We clear the cycle by allocating o_3 to a_2 and o_2 to a_3 and remove the agents. Since no agents remain, the algorithm terminates with the same assignment as that made by SD.

The above matrices are doubly stochastic, as expected: each column and each row has sum 1. The matrix representing the rank distribution arising from using SD with picking order a_1, a_2, a_3 and TTC with initial endowment o_i to agent a_i is $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$ since agents a_1 and a_3 receive their top object, and agent a_2 receives their second-top object. The rank distribution matrix for RSD is $\begin{bmatrix} \frac{3}{6} & \frac{1}{6} & \frac{2}{6} \\ \frac{3}{6} & \frac{3}{6} & 0 \\ \frac{5}{6} & 0 & \frac{1}{6} \end{bmatrix}$.

In Example 3.2 the rows of the rank distribution matrices are very far from being equal. The case where they are equal gives the main new concept of this paper.

Definition 3.3. Let P be a probability measure on the set of all profiles for a given n . The mechanism $\mathcal{A}$ is **order-symmetric with respect to P** if for all j , the quantity

$$E_P[D_{\pi, \mathcal{A}}(r, j)] := \sum_{\pi \in \Pi(A, O)} P(\pi) D_{\pi, \mathcal{A}}(r, j)$$

is independent of r .

Thus a mechanism is order-symmetric if and only if in the expectation of the rank distribution matrix, all rows are equal. The following example illustrates the idea. **Impartial Culture** (IC) is the probability measure obtained when each agent independently samples from the uniform measure on all permutations of the objects. This is just the uniform measure on profiles.

Example 3.4. Consider the case of agents a_1, a_2, a_3 and objects o_1, o_2, o_3 as in Example 3.2, and assignment mechanism SD with picking order a_1, a_2, a_3 . By symmetry we may assume that the preference order of a_1 is $o_1 \succ o_2 \succ o_3$. This leaves 36 possibilities for the other two preference orders. Then a_1 always chooses o_1 , a_2 takes whichever of o_2 and o_3 is preferred and a_3 is left with the last object, which is equally likely to take any rank in her preference order. Unless the first choice of a_2 is o_1 , he gets his first object. The expected rank distribution matrix under IC is therefore $\begin{bmatrix} \frac{1}{3} & 0 & 0 \\ \frac{2}{3} & \frac{1}{3} & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$.

For TTC with initial endowment that assigns o_i to a_i for $i \in \{1, 2, 3\}$ and for RSD, the expected rank distribution matrix under IC follows from computing the outcome on each profile and is $\begin{bmatrix} \frac{2}{3} & \frac{2}{9} & \frac{1}{9} \\ \frac{2}{3} & \frac{2}{9} & \frac{1}{9} \\ \frac{2}{3} & \frac{2}{9} & \frac{1}{9} \end{bmatrix}$.

From Example 3.4, we see that RSD and TTC are order symmetric with respect to IC for $n = 3$, but not with respect to the point mass concentrated at the profile from Example 3.2. Also, SD with picking order a_1, a_2, a_3 is not order symmetric with respect to either measure. This is not surprising; the first chooser does strictly better than the second when both have the same first preference, and just as well otherwise. We will soon see that this example generalizes, with **TTC-RE** offering stronger ex post order symmetry guarantees than **RSD**.

Before proceeding, we observe that order symmetry is a property of a mechanism that cannot be verified by examining individual allocations in isolation. In this respect it is similar to anonymity (we formalize a connection in the next section), but of a different flavor to guarantees such as equal treatment of equals or envy-freeness, both of which define behavior that a mechanism must satisfy on each and every input independently.

3.2 Relation to anonymity

When order symmetry is combined with sufficient symmetry in the probability measure, we obtain a property that is a relaxation of anonymity. Recall that $G = \text{Sym}(A)$ denotes the group of all agent permutations.

Definition 3.5. An **anonymous probability measure** on Π is one for which $P(g \cdot \pi) = P(\pi)$ for all $g \in G$ and all $\pi \in \Pi$.

Note that a probability measure is anonymous if and only if it takes a constant value on each orbit of G . Further, every measure that is uniform on a subset of profiles that is a union of G -orbits is anonymous. Examples of anonymous measures include Impartial Culture (where each profile has the same weight), Impartial Anonymous Culture (IAC, where any two profiles within the same orbit have the same weight), and Impartial Anonymous Neutral Culture (IANC) (Egencioglu and Giritligil 2013). Other examples include any measure concentrated on a single profile in which all agents have the same preference, the uniform measure on single-peaked profiles, and indeed the uniform measure on any set of profiles defined without reference to an order of agents.

Theorem 3.6. *If $\mathcal{A}$ is an anonymous assignment mechanism and P is an anonymous probability measure on $\Pi(\mathcal{A}, \mathcal{O})$ then $\mathcal{A}$ satisfies order symmetry with respect to P .*

The proof is presented in Appendix B. It relies on the fact that each permutation g of the agents permutes the rank distribution matrix in the same way (since the mechanism is anonymous), and that each profile in the orbit of g is equally likely (by anonymity of the probability measure).

In the following section, we use order symmetry as a criterion to discriminate between assignment mechanisms on ex post fairness grounds.

4 Satisfying Order Symmetry

We focus our attention in this work on order symmetry with respect to anonymous probability measures (see Appendix J for results regarding more general probability measures). By Theorem 3.6, in the presence of such a measure, order symmetry is a fairness property that is implied by anonymity. It is not a priori obvious whether order symmetry can be satisfied by discrete mechanisms (we know that anonymity cannot). As the following example shows, no discrete mechanism can satisfy order symmetry with respect to *all* anonymous measures.

Example 4.1. *Consider a profile where all agents have the same preference order, and let P put all its weight on this profile. Note that P is anonymous. Then every discrete assignment mechanism must allocate some agent her most-preferred object, while other agents will not get their first choice, so that the mechanism cannot be order symmetric with respect to P .*

On the other hand, some probability measures make it easy to be order symmetric. Consider a measure with all its weight on the set of profiles for which all agents have a different first preference. Every Pareto efficient discrete assignment mechanism will assign each agent to their top choice, and thus order symmetry is satisfied.

4.1 Fully symmetric probability measures

In the light of Example 4.1, we will now further restrict attention to a natural subclass of anonymous measures: those which are additionally invariant to permutations of the objects.

Definition 4.2. *Let $H = \text{Sym}(\mathcal{O})$ be the group of all permutations of the set of objects. Note that H acts naturally on $\Pi(\mathcal{A}, \mathcal{O})$ by $(h \star \pi)(a) = (h(\pi(a)_1), \dots, h(\pi(a)_n))$ and also on assignments by $(h \star \alpha)(a) = h(\alpha(a))$. A **neutral probability measure** on $\Pi(\mathcal{A}, \mathcal{O})$ is one for which $P(h \star \pi) = P(\pi)$ for all $h \in H$ and all $\pi \in \Pi$.*

Example 4.3. *Consider the set of all profiles for which the last $n-1$ agents all have the same preference order, and agent 1 always has a different one. If $n \geq 3$ then the uniform measure on this set is neutral but not anonymous.*

The actions of G and H commute with each other (reordering agents, then objects, gives the same result as reordering objects, then agents), so that

$$g \cdot (h \star \pi) = h \star (g \cdot \pi) \quad \text{for all } g \in G, h \in H.$$

We have given a definition that allows for more general sets of agents and objects. In the specific case of the house allocation problem where $|\mathcal{A}| = |\mathcal{O}| = n$ and we choose a fixed order on $\mathcal{A}$ and $\mathcal{O}$ as we have done above, we may identify both $\mathcal{A}$ and $\mathcal{O}$ with $\{1, \dots, n\}$ and G and H with the symmetric group S_n . We can now define the class of fully symmetric probability measures.

Definition 4.4. *A probability measure is **fully symmetric** if and only if $P(g \cdot (h \star \pi)) = P(\pi)$ for all $g \in G, h \in H$.*

Proposition 4.5. *For a probability measure P on Π , the following properties are equivalent.*

- (i) *P is fully symmetric.*
- (ii) *P is anonymous and neutral.*
- (iii) *P takes a constant value on each orbit of $G \times H$*

Examples of fully symmetric measures include IC, IAC, IANC and indeed the uniform measure on any set of profiles defined without reference to an order of agents or objects. A measure concentrated on a single profile in which all agents have the same preference is anonymous, but not fully symmetric since it is not invariant to permuting the objects.

4.2 TTC is order symmetric

We now show our first main result, that TTC is order symmetric with respect to any fully symmetric probability measure. The result generalizes Example 3.4, which showed that TTC is order symmetric with respect to IC for $n = 3$. Intuitively, the result is true because, under the assumption that each agent and object is symmetric, no agent is systematically advantaged by their initial endowment. The proof, along with an illustrative example, appears in Appendix D.

Theorem 4.6 (TTC is order symmetric). *TTC with any fixed endowment is order symmetric with respect to every fully symmetric probability measure.*

Because we know that TTC with fixed endowment is strategyproof and Pareto efficient, we obtain the following corollary.

Corollary 4.7. *TTC-RE is ex ante anonymous, ex post strategyproof, ex post Pareto efficient, and ex post order symmetric with respect to every fully symmetric probability measure.*

Since we know that **TTC-RE** is an implementation of RSD, we can state Corollary 4.7 in the following way.

Corollary 4.8. *RSD can be implemented as a lottery over discrete assignment mechanisms each of which is strategyproof, Pareto efficient and order symmetric with respect to every fully symmetric probability measure.*

Note that the standard implementation of RSD (that is, **RSD**) violates order symmetry not only with respect to Impartial Culture, as we have seen in Example 3.4, but also with respect to all but a very special set of measures.

Theorem 4.9. *RSD violates ex post order symmetry with respect to any probability measure that assigns positive probability to a profile in which two agents have the same most-preferred object.*

Proof. Consider SD with a fixed picking order. The first agent in the picking order always receives their most-preferred object. Thus, the only way for all rows of the expected rank distribution matrix to be equal is if all agents always receive their most-preferred object. However, on any profile where two agents have the same most-preferred object, one of them will receive a lower-ranked object, violating order symmetry. $\square$

In summary, **RSD** and **TTC-RE** induce the same fractional assignment mechanism and therefore necessarily share any ex ante properties. They are also ex post Pareto efficient and ex post strategyproof. However there is a substantial difference: only **TTC-RE** is ex post order symmetric. The usual implementation of RSD does not give the correct type of lottery over discrete assignment mechanisms, and implementing it using **TTC-RE** fixes this problem.

4.3 Order symmetry and efficiency

We have seen that **TTC-RE** satisfies many desirable properties. However, it is not ex ante ordinally efficient (Bogomolnaia and Moulin 2001). This begs a natural question: does there exist a randomized assignment mechanism that is ex ante ordinally efficient, ex ante anonymous, and ex post order symmetric?

To answer this question, we define a discrete assignment mechanism that we call the **Rank-Maximal Matching** (RMM) mechanism. A rank-maximal matching is an assignment that maximizes the number of agents who receive their most-preferred object, and subject to that maximizes the number who receive their second most-preferred object, and so on (Irving 2003, Irving et al. 2006). RMM with a fixed tiebreak order ρ first finds the complete set of rank-maximal matchings, and then breaks ties by having each agent in turn (according to ρ) find their most preferred object of those she gets in one of the remaining assignments, and eliminates all assignments that give her anything worse. Clearly this terminates with a single assignment, since every agent gets exactly one object. This mechanism is exactly the welfare maximization method proposed by Featherstone (2020), since a rank-maximal matching is one that maximizes social welfare for agents with lexicographic utilities.

For any fixed tiebreak order ρ , RMM is clearly not order symmetric, systematically favoring agents higher in the tiebreak order. However, we propose breaking ties according to $\pi(i)$, the preference list of agent i (where each object o_j is associated with agent a_j). This tiebreaking method harnesses the randomness from the input to avoid biasing the mechanism in favor of any particular agent, provided that the input distribution is sufficiently symmetric.

Example 4.10. Consider the profile where agents a_1, a_2, a_3 have respective preferences over objects o_1, o_2, o_3 as follows.

$$\begin{aligned} a_1 : o_3 &\succ o_1 > o_2 \\ a_2 : o_3 &\succ o_1 > o_2 \\ a_3 : o_2 &\succ o_3 > o_1 \end{aligned}$$

The two rank maximal matchings are $a_1 : o_3, a_2 : o_1, a_3 : o_2$ and $a_1 : o_1, a_2 : o_3, a_3 : o_2$, since both give two agents their most-preferred object and one agent their second most-preferred.

Suppose the tiebreak order is given by $\pi(1) = 3, 1, 2$. Since agent 3 is indifferent between the two assignments, neither is eliminated before agent 1's turn. Agent 1 chooses the assignment she prefers of the two, yielding assignment matrix $\begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Lemma 4.11. Fix an agent i . Then RMM with tiebreaking order $\pi(i)$ is order symmetric with respect to any fully symmetric probability measure.

The proof of Lemma 4.11 appears in Appendix E. Observe that the same trick can be applied to serial dictatorship as well, by tying the picking order to the preference of some fixed agent. However, doing so achieves order symmetry at the expense of strategyproofness, while we can achieve both properties simultaneously by using TTC.

Equipped with an order symmetric discrete assignment mechanism, defining a randomized mechanism that is ex ante anonymous while preserving ex post order symmetry is straightforward, as per the construction in Appendix H.2. In this case, it is equivalent to simply choosing the agent whose preference defines the tiebreak order uniformly at random. We term the resulting randomized assignment mechanism **Rank-Maximal Matching with Random Agent tiebreaking** (RMM-RA). As well as being ex post order symmetric, **RMM-RA** is ex ante ordinally efficient.

Theorem 4.12. **RMM-RA** is ex ante ordinally efficient, ex ante anonymous, and ex post order symmetric with respect to every fully symmetric measure.

Since we have shown order symmetry in Lemma 4.11, it only remains to show efficiency, which follows easily from the definition of a rank maximal matching. The proof appears in Appendix F. As it turns out, the proof implies that RMM-RA actually satisfies the stronger efficiency property of ex ante rank efficiency (Featherstone 2020).

Despite its strong efficiency guarantee, **RMM-RA** does have significant drawbacks. It necessarily violates ex ante strategyproofness, since ex ante anonymity and ex ante ordinal efficiency are known to be incompatible with ex ante strategyproofness (Zhou 1990, Bogomolnaia and Moulin 2001). Additionally, computing all rank-maximal matchings is a #P-hard computational problem (Ghosal et al. 2014), rendering the mechanism impractical for most real-world problems. However, it might be possible that

rather than computing rank-maximal matchings and then breaking ties among them, some other set of assignments could be computed (in polynomial time) that still guarantee ordinal efficiency. In this case, the same tiebreaking procedure could be used to guarantee ex post order symmetry as well.

Let us also mention the *Probabilistic Serial* (PS) mechanism of Bogomolnaia and Moulin (2001). PS constructs an assignment via a “simultaneous eating” procedure: agents all eat objects at unit speed, switching down their preference order to the next best remaining once an object is completely eaten. As a fractional assignment mechanism, PS is ordinally efficient, anonymous, and polynomial-time computable. However, its definition does not lend itself to a natural implementation in the same way as, say, RSD does. Of course, the method from Appendix G can be applied to exhibit PS as a lottery over discrete assignment mechanisms, but we have no guarantee that such an implementation would be ex post order symmetric. If such an implementation of PS does exist (and can be computed efficiently), then it would be a natural candidate for practical use, especially given its status as an existing mechanism with a somewhat more natural description than **RMM-RA**. Resolving this question is a compelling problem that we leave open.

To conclude this section, we briefly contextualize our theoretical results. It is known that ex ante ordinal efficiency, ex ante strategyproofness, and ex ante anonymity are incompatible (Bogomolnaia and Moulin 2001). The literature suggests two common approaches in light of this difficulty. The first, typically achieved by implementing random serial dictatorship, is to sacrifice ex ante ordinal efficiency, resorting instead to ex post efficiency. The second is to sacrifice strategyproofness by using the probabilistic serial mechanism. In both cases, we have shown that ex post order symmetry (with respect to a fully symmetric probability measure) can be additionally achieved at no cost, via the use of TTC and Rank Maximal Matching, respectively.

5 Discussion

We have defined and analysed order symmetry, a novel and natural fairness concept that requires all agents to do equally well in expectation. When associated with a sufficiently symmetric probability measure on profiles, order symmetry is a relaxation of anonymity that, unlike anonymity, can be satisfied by discrete algorithms. For other measures, order symmetry cannot be satisfied exactly, and some order bias is inevitable. In situations where order bias is a consideration, our analysis suggests that the standard implementation of **RSD** should be avoided, and **TTC** used instead. The latter achieves the same ex ante results as the former, while being much fairer ex post, in both theory and simulation. In Appendix A we explore several common algorithms (random serial dictatorship, top trading cycles, and two variants of the Boston mechanism) in terms of order bias, a measure that quantifies violations of order symmetry. Of these standard algorithms, the Boston mechanisms (under truthful behavior) perform well in terms of order bias, never satisfying order symmetry exactly but significantly outperforming even **TTC** on synthetic Mallows data for all but very large values of dispersion parameter. Our results suggest that Naive Boston in particular should be considered for a designer wishing to minimize the influence of the mechanism’s randomness and for whom strategyproofness is less important than fairness.

We have additionally shown that ex ante ordinal efficiency can be achieved in conjunction with ex post order symmetry with respect to any fully symmetric probability measure. While the **RMM-RA** mechanism that we invent is clearly not feasible for practical implementation—it is computationally intractable and employs an unintuitive tiebreaking procedure—we see it as a valuable proof of concept that illustrates the compatibility of these compelling properties.

We conclude with some additional observations and directions for future work.

Order Symmetric Implementations. As we have already mentioned, it would be particularly interesting to determine whether or not it is possible to implement the Probabilistic Serial mechanism, known to be ex ante ordinally efficient and ex ante weakly strategyproof (Bogomolnaia and Moulin 2001), in a way that is ex post order symmetric. An assignment mechanism is weakly strategyproof if an agent cannot obtain an assignment by misreporting that stochastically dominates their assignment from truth-telling.

More generally, can we determine whether a given fractional assignment mechanism can be implemented over discrete assignment mechanisms satisfying order symmetry? It is known that the analogous question is computationally hard for some other properties. For example, determining whether a randomized mechanism is ex-post efficient, and hence can be expressed as a lottery over Pareto optimal assignment mechanisms, is NP-complete (Aziz et al. 2015).

Multi-unit assignment. We have not considered the case where the number of objects is more than the number of agents and a bundle of objects is assigned to each agent, but the idea of order bias remains relevant there and our definition could be generalized by extending preferences over individual objects to preferences over sets of objects (Barberà et al. 2004). Indeed many common mechanisms, including sports league drafts, utilize a “picking sequence” (order in which agents choose objects, one by one) and this sequence has long been known to influence outcomes substantially. Many attempts have been made to balance outcomes between agent positions (Kohler and Chandrasekaran 1971, Brams and Taylor 2000, Bouveret and Lang 2011, Beynier et al. 2019, Kalinowski et al. 2013), and order bias may help to compare different picking sequences.

We note that there is already a substantial stream of work aimed at measuring fairness in this setting. Concepts such as envy-freeness up to one object (Lipton et al. 2004), envy-freeness up to any object (Caragiannis et al. 2019), and proportionality up to one object (Conitzer et al. 2017) have been widely studied, and prior work has considered ex post fairness in conjunction with ex ante guarantees in more general multi-object assignment settings (Budish et al. 2013, Akbarpour and Nikzad 2020, Freeman et al. 2020, Aziz 2020, Babaioff et al. 2021). None of this work is relevant to the house allocation setting that we consider in this work, since the fairness guarantees are too permissive when each agent receives only a single object.

Two-sided matching. One may wonder why we have focused on one-sided matching and not the two-sided case. The answer is that, from the perspective of order symmetry, the two-sided case is significantly less interesting! In the two-sided setting, agents on both sides of the market express a preference, in contrast to the one-sided setting in which the objects do not have preferences over the agents. Two-sided matching mechanisms are therefore able to leverage the preferences of one side as a source of asymmetry over the other side, rather than having to introduce that asymmetry as part of the mechanism (as a tiebreak order or initial endowment). Two-sided mechanisms such as deferred acceptance and (the two-sided version of) TTC are defined without any reference to an ordering of agents and therefore satisfy a natural generalization of anonymity without any need for randomization.

General probability measures. A compelling direction for future work is to dive more deeply into the possibilities and limitations for general probability measures. While measures such as IC are a useful baseline, they are rarely realistic, and our empirical investigation suggests that the choice of probability measure has a large influence on order bias. For practical application, we imagine order symmetry being incorporated into mechanism design in a two step process. First, the designer estimates (say, from historical data) the probability measure from which the preference profile will be drawn. Second, the designer can search for a mechanism with low order bias with respect to this measure.

Other order bias definitions. The way we have defined order bias is not the only possibility. Abstractly, we could use any matrix norm to measure the distance between the given rank distribution matrix and the space of rank one matrices. It is unclear whether this is more compelling axiomatically than the approach we have taken, but it is worth further exploration.

Obvious strategyproofness. We have used the concept of order symmetry to distinguish between **RSD** and **TTC-RE**. Another distinction between the two can be made via the notion of obvious strategyproofness (Li 2017). When implemented via the usual sequential choosing procedure, SD is obviously strategyproof, while TTC has no obviously strategyproof implementation. In the language of this paper, RSD is ex post obviously strategyproof, while TTC with random endowment is not. Given that obvious strategyproofness increasingly appears to be a desirable property for practical application (Schmelzer 2017, Troyan 2019), it would be natural to consider the possibility of achieving it in conjunction with order symmetry.

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998. URL <http://www.jstor.org/stable/2998580>.
- Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

- David J Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*, pages 3–15. Springer, 2004.
- Mohammad Akbarpour and Afshin Nikzad. Approximate random allocation mechanisms. *Review of Economic Studies*, 87(6):2473–2510, 2020.
- Haris Aziz. Simultaneously achieving ex-ante and ex-post fairness. In *International Conference on Web and Internet Economics*, pages 341–355. Springer, 2020.
- Haris Aziz, Simon Mackenzie, Lirong Xia, and Chun Ye. Ex post efficiency of random assignments. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1639–1640. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. Best-of-both-worlds fair-share allocations. *arXiv preprint arXiv:2102.04909*, 2021.
- Sophie Bade. Random serial dictatorship: the one and only. *Mathematics of Operations Research*, 45(1): 353–368, 2020.
- Salvador Barberà, Walter Bossert, and Prasanta K Pattanaik. Ranking sets of objects. In *Handbook of utility theory*, pages 893–977. Springer, 2004.
- Aurélien Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams. Efficiency, sequenceability and deal-optimality in fair division of indivisible goods. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 900–908, 2019.
- Garrett Birkhoff. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumán Rev, Ser. A, no. 5*, pages 147–151, 1946.
- Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001. URL <http://www.sciencedirect.com/science/article/pii/S0022053100927108>.
- Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Twenty-Second International Joint Conference on Artificial Intelligence*. Citeseer, 2011.
- Steven J. Brams and Alan D. Taylor. *The win-win solution: Guaranteeing fair shares to everybody*. WW Norton & Company, 2000.
- Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- Yan Chen and Tayfun Sönmez. Improving efficiency of on-campus housing: An experimental study. *American Economic Review*, 92(5):1669–1686, 2002.
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 629–646, 2017.
- Hervé Crès and Hervé Moulin. Scheduling with opting out: Improving upon random priority. *Operations Research*, 49(4):565–577, 2001.
- Xiaotie Deng, Yansong Gao, and Jie Zhang. Smoothed and average-case approximation ratios of mechanisms: Beyond the worst-case analysis. *arXiv preprint arXiv:1705.07200*, 2017.
- Ömer Eğecioğlu and Ayça E. Giritligil. The impartial, anonymous, and neutral culture model: a probability model for sampling public preference structures. *Journal of Mathematical Sociology*, 37(4):203–222, 2013.
- Haluk I. Ergin. Consistency in house allocation problems. *Journal of Mathematical Economics*, 34(1):77–97, 2000.
- Clayton Featherstone. Rank efficiency: Modeling a common policymaker objective. 2020.
- Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *International Symposium on Algorithmic Game Theory*, pages 1–12. Springer, 2014.
- D. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–98, 1967.
- Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 21–22, 2020.
- Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.
- Yansong Gao and Jie Zhang. Average-case analysis of the assignment problem with independent preferences. *arXiv preprint arXiv:1906.00182*, 2019.

- Pratik Ghosal, Meghana Nasre, and Prajakta Nimbhorkar. Rank-maximal matchings—structure and algorithms. In *International Symposium on Algorithms and Computation*, pages 593–605. Springer, 2014.
- Patrick Harless and Vikram Manjunath. Learning matters: Reappraising object allocation rules when agents strategically investigate. *International Economic Review*, 59(2):557–592, 2018.
- Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, pages 293–314, 1979. URL <http://www.jstor.org/stable/1832088>.
- Robert W Irving. Greedy matchings. *Technical Report Tr-2003-136, University of Glasgow*, 2003.
- Robert W Irving, Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E Paluch. Rank-maximal matchings. *ACM Transactions on Algorithms (TALG)*, 2(4):602–610, 2006.
- Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 227–233, 2013.
- David A. Kohler and R. Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.
- Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. Size versus truthfulness in the house allocation problem. *Algorithmica*, 81(9):3422–3463, 2019.
- Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017.
- Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.
- Xinghua Long. *Essays on house allocation problem*. PhD thesis, Department of Economics, Texas A&M University., 2016. URL <https://hdl.handle.net/1969.1/157810>.
- Xinghua Long and Rodrigo A Velez. Balanced house allocation. *arXiv preprint arXiv:2109.01992*, 2021.
- Colin L. Mallows. Non-null ranking models. I. *Biometrika*, 44(1/2):114–130, 1957.
- Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. Fair dynamic rationing. *Available at SSRN 3775895*, 2021.
- Timo Mennle and Sven Seuken. The Naive versus the Adaptive Boston Mechanism. *arXiv preprint arXiv:1406.3327*, 2014. URL <http://arxiv.org/abs/1406.3327>.
- Timo Mennle and Sven Seuken. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191:105144, 2021.
- Alexander S. Nesterov. Fairness and efficiency in strategy-proof object allocation mechanisms. *Journal of Economic Theory*, 170:145–168, 2017.
- Parag A Pathak and Jay Sethuraman. Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6(1):1–17, 2011.
- Geoffrey Pritchard and Mark C. Wilson. Asymptotic welfare performance of boston assignment algorithms. 2022. In preparation.
- Marek Pycia and M Utku Ünver. Incentive compatible allocation and exchange of discrete resources. 2014. Working paper.
- Marek Pycia and M Utku Ünver. Decomposing random mechanisms. *Journal of Mathematical Economics*, 61:21–33, 2015.
- Alvin E. Roth. Incentive compatibility in a market with indivisible goods. *Economics letters*, 9(2):127–132, 1982.
- Alvin E. Roth and Andrew Postlewaite. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4(2):131–137, 1977.
- Daniela Saban and Jay Sethuraman. House allocation with indifferences: a generalization and a unified view. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 803–820, 2013.
- Daniela Saban and Jay Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.
- André Schmelzer. Strategy-proofness of stochastic assignment mechanisms. *MPI Collective Goods Preprint*, 2017(13), 2017.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1): 23–37, 1974. URL <http://www.sciencedirect.com/science/article/pii/0304406874900330>.
- Tayfun Sönmez and M. Utku Ünver. House allocation with existing tenants: A characterization. *Games and Economic Behavior*, 69(2):425–445, 2010.
- William Thomson. Children crying at birthday parties. Why? *Economic Theory*, 31(3):501–521, 2007.
- Peter Troyan. Obviously strategy-proof implementation of top trading cycles. *International Economic Review*, 60(3):1249–1261, 2019.

- J. von Neumann. A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem. In W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 2, pages 5–12. Princeton University Press, 1953.
- Lin Zhou. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1): 123–135, 1990.

A Violations of order symmetry

By Corollary 4.7 and Theorem 4.12, there are randomized assignment mechanisms with strong axiomatic properties including order symmetry. However it is still important for us to measure the deviation from order symmetry, for several reasons. First, some preference distributions may not admit any order symmetric mechanism, and we may wish to get as close as we can. Second, a small failure of order symmetry may be tolerable in practice, if for example it leads to higher overall welfare. Third, as in the example of sports drafts, sometimes we do not want order symmetry, but rather a bias toward some agent(s). We therefore investigate how to measure the failure of order symmetry to hold. We first define a notion of order bias and examine some of its theoretical consequences, then examine the performance of four common rules empirically.

A.1 Order bias

To define a quantitative measure of bias in the absence of knowledge of preferences beyond ordinal information, we force a common utility function on agents via a scoring rule.

Definition A.1. A *positional scoring rule* is given by a sequence s of real numbers $s_1 \geq s_2 \geq \dots \geq s_n$ with $s_1 > s_n$.

Commonly used scoring rules include **plurality** defined by $(1, 0, 0, \dots, 0)$, **antiplurality** defined by $(1, 1, 1, \dots, 0)$ and **Borda** defined by $s_i = (n - i)/(n - 1)$. By equating the entries in a scoring rule with utilities, we can define the expected utility of an agent for some measure on profiles.

Definition A.2. The *expected utility* of agent r with respect to a scoring rule s , mechanism $\mathcal{A}$ and measure P on profiles is

$$U(r) := U_{\mathcal{A}, s, P}(r) = \sum_j s_j E_P[D_{\pi, \mathcal{A}}(r, j)].$$

In other words, U is a vector formed by multiplying the expected rank distribution matrix by the fixed utility vector. We now require a measure of bias. An obvious choice is to consider the maximum difference between two agents' expected utilities, normalized by the difference in utility for receiving the most- and least-preferred objects:

$$\beta_n(\mathcal{A}; s; P) = \frac{\max_{1 \leq p, q \leq n} |U(p) - U(q)|}{s_1 - s_n}.$$

With the given normalization, $0 \leq \beta_n(\mathcal{A}; s; P) \leq 1$, with order symmetry implying $\beta_n(\mathcal{A}; s; P) = 0$, while the equality $\beta_n(\mathcal{A}; s; P) = 1$ is attained only if there are fixed agents A, B such that with probability 1, A attains its first choice and B attains its worst choice.

The notion of order bias allows us to make formal comparisons regarding the average case (un)fairness of different rules. As a simple case, consider the impartial culture measure on profiles and the serial dictatorship mechanism with respect to a fixed picking order (we already know that TTC with fixed endowment has zero order bias).

Theorem A.3. Let $\mathcal{A}$ be SD with picking order equal to $\{a_1, \dots, a_n\}$, and let P be IC. Then

$$E_P[D_{\pi, \mathcal{A}}(r, j)] = \begin{cases} \frac{\binom{n-j}{r-j}}{\binom{n}{r-1}} & \text{if } j \leq r \\ 0 & \text{if } j > r. \end{cases}$$

Proof. By the time agent r gets an object, a random subset S of $r - 1$ of the n objects is already taken. The probability that his j th choice is the best object left is the probability that S includes his first $j - 1$ choices, but not the j th choice. The number of subsets satisfying this constraint is $\binom{n-j}{r-j}$ because we have to choose the remaining $r - j$ objects from $n - j$ possibilities. $\square$

In particular if $r = n$, then agent r is equally likely to get each possible object.

Proposition A.4. *We have $\beta_n(SD; s; IC) = \frac{s_1 - \bar{s}}{s_1 - s_n}$, where $\bar{s}$ is the mean of s .*

Proof. Agent 1 always attains its first choice under SD. The last agent receives each object with equal probability. $\square$

For example, for the Borda rule $\beta_n(SD; s; IC) = 1/2$ while for plurality $\beta_n(SD; s; IC) = 1 - 1/n$, and for antiplurality $\beta_n(SD; s; IC) = 1/n$.

We conclude this section by briefly discussing the relationship between welfare and order bias.

Definition A.5. *Let n be a positive integer, $\mathcal{A}$ an assignment mechanism, s a scoring rule and P a distribution over profiles with n agents. The **utilitarian welfare** $W_n(\mathcal{A}; s; P)$ for $\mathcal{A}$ with respect to s and P is the expectation under P of the arithmetic mean score received by an agent, while the **egalitarian welfare** $E_n(\mathcal{A}; s; P)$ is the expectation of the minimum score. The **Nash welfare** $N_n(\mathcal{A}; s; P)$ is the expectation of the geometric mean of agent scores.*

From the definitions it follows immediately that for a given $(\mathcal{A}; s; P)$, $E_n \leq N_n \leq W_n$ for all n .

Proposition A.6. *For each $n, s, P, \mathcal{A}$, we have $\beta_n(\mathcal{A}; s; P) + E_n(\mathcal{A}; s; P) \leq 1$.*

Proof. Let M and m denote the maximum and minimum over agents of the normalized expected utility of an agent. Then $m \geq E_n(\mathcal{A}; s; P)$ while $\beta_n(\mathcal{A}; s; P) = M - m$, so the result follows since $M \leq 1$. $\square$

Thus low order bias is a necessary condition for high egalitarian welfare. It is certainly not sufficient, since for example we could run **TTC** after reversing the preference order of each agent.

If order symmetry holds, then the expected welfare of each agent is the same, and hence equal to the utilitarian welfare. However the relationship between order bias and utilitarian (or Nash) welfare is not clear in general.

A.2 Empirical results

In this section we use simulation to study the order bias of four well known discrete assignment algorithms: TTC with fixed endowment, SD with fixed picking order, Naive Boston (NB) with fixed tiebreak order and Adaptive Boston (AB) with fixed tiebreak order (see Appendix L for a discussion of alternative tiebreaking mechanisms in the Boston mechanisms).

The **Boston mechanism** for school choice proceeds in rounds. At round i , all unmatched agents are asked to submit their i th choice, and they are allocated that object unless there is a conflict, in which case a fixed order ρ is used as a tiebreak to allocate it to one of them. Alternatively, agents submit their first preference and if it has already been taken, they go back to the end of the queue of agents (initially built from ρ) and wait to submit their 2nd choice, etc. In the present situation we consider the restriction to the housing allocation model, so that each “school” is a single object.

Mennle and Seuken (2014) discuss what they call the **Adaptive Boston** mechanism (they call the original variant **Naive Boston**). In this case, in round i , all remaining agents submit a bid for their highest remaining object, rather than for the highest object for which they have not yet bid, as in the naive version. It is well known that the both variants are Pareto efficient but violate strategyproofness (Mennle and Seuken 2014, Abdulkadiroğlu and Sönmez 2003), although the adaptive version does satisfy a weaker incentive property known as partial strategyproofness (Mennle and Seuken 2014, 2021).

For the remainder of this section, we will use the abbreviations SD, TTC, NB, and AB not to represent fractional assignment mechanisms, but to refer to the discrete form of the mechanisms when endowed with some fixed endowment/tiebreak order.

It is clear from their definitions that, as for SD, the Boston mechanisms are order symmetric only for very particular probability measures. Most of the time, an agent at position i in the tiebreak order will obtain strictly higher expected utility than an agent at position $j > i$ in the order, since agent i only has to compete with agents $\{1, \dots, i-1\}$ while agent j competes with $\{1, \dots, j-1\}$. This is very clearly shown in the numerical results that follow.




Figure 1: (left) Borda order bias, Impartial Culture, mean of 1000 simulations, (right) distribution of rank received for last agent, $n = 16$, Impartial Culture, mean of 1000 simulations.

Impartial Culture

Figure 1 explores the order bias of all four algorithms. On the left, we plot the sample mean under IC (from 1000 independent samples for each $n \in \{2, 4, 8, 16, 32\}$) of order bias with respect to the Borda utility (in Appendix L we examine other utility functions and observe similar results to those presented in this section). In line with the theory, the value for TTC is near zero and for SD near $1/2$. The Boston mechanisms exhibit intermediate performance, improving as n increases, with the naive version outperforming its adaptive counterpart with respect to order bias. (This phenomenon persists throughout our empirical results. Providing a complete theoretical justification remains an open question, but we give a partial explanation in Appendix L.)

The Borda order bias reflects the difference in the *average* rank of the object received by the agents with highest and lowest average Borda utility, but not the distribution of ranks. On the right side of Figure 1, we plot the distribution of ranks for the agent with lowest average Borda utility. Recall that for all algorithms other than TTC, the highest agent in the priority order always gets its top choice.

Of course, there are design criteria other than order bias. Our results in Appendix L.4 show that the differences in welfare performance of the algorithms are very small compared to their differences in order bias. In particular, as expected due to the results of Abdulkadiroğlu and Sönmez (1998) and Bade (2020), SD and TTC are almost identical in welfare.

Mallows model

Impartial culture is a helpful and tractable baseline, but often does not reflect realistic preference structures. In particular, in real-world assignment problems, some objects are typically more highly sought after than others but this is not captured by Impartial Culture. To explore a more realistic preference model, we consider the Mallows model (Mallows 1957). The Mallows model is parameterized by a **reference order** σ and a **dispersion parameter** $\phi \in (0, 1]$. Let $\succ$ be a strict linear order. Then the Mallows model specifies that the probability an agent has preference order $\succ$ is $P(\succ) = P(\succ | \sigma, \phi) = \frac{1}{Z} \phi^{d(\succ, \sigma)}$, where Z is a normalization parameter and $d(\succ, \sigma)$ is the number of pairwise disagreements between $\succ$ and σ (that is, the number of pairs of objects that are ordered differently by $\succ$ than by σ), known as the Kendall-Tau distance. To generate preference profiles, we sample the preference of each agent independently from the same Mallows distribution. Note that the Mallows model is unimodal at σ , with lower ϕ implying a more concentrated distribution of preferences around σ . Note that when $\phi = 1$ the distribution is maximally dispersed, yielding Impartial Culture. Unless $\phi = 1$ the Mallows measure is not fully symmetric, since some orderings have higher probability than others in the same H -orbit.

Figure 2 is analogous to Figure 1, but with a Mallows distribution with $\phi = 0.6$. In this case, we see a consistent advantage for TTC over SD, but less than for IC. For intuition as to why, when some objects are systematically preferred to others then TTC is biased in favor of agents who receive favorable objects in the preliminary assignment. In the limit as the dispersion parameter goes to zero (all agents have the same preference), being initially assigned the k -th most favorable object is exactly equivalent to choosing k -th in the SD order. The Boston mechanisms, especially the naive version, achieve significantly lower order bias still. The drivers of this pattern are clear from the right hand side of the figure. Under SD and TTC, the worst-off agent rarely receives their most-preferred object. However, the Boston mechanisms




Figure 2: (left) Borda order bias, Mallows ($\phi = 0.6$), mean of 1000 simulations, (right) distribution of rank received for last agent, $n = 16$, Mallows ($\phi = 0.6$), mean of 1000 simulations.

(by design) give all agents a high probability of receiving their most-preferred object. Remarkably, under Naive Boston, the probability of receiving an object is (approximately) decreasing in the rank of that object, even for the last agent in the tiebreak order, which accounts for the low order bias.




Figure 3: (left) Borda order bias, $n = 32$, mean of 1000 simulations, (center) Borda utilitarian welfare, $n = 32$, mean of 1000 simulations, (right) Borda Nash welfare, $n = 16$, mean of 1000 simulations.

Finally, Figure 3 explores the relationship between the Mallows dispersion parameter and order bias, utilitarian welfare, and Nash welfare. Most strikingly, on the left, we see that the Boston mechanisms have lower order bias than TTC except for very high values of dispersion parameter. This suggests that the Boston mechanisms, and Naive Boston in particular, may be a better choice than TTC for minimizing order bias when some objects are systematically preferred over others.

B Proof of Theorem 3.6

Suppose that $\mathcal{A}$ is anonymous, and consider agent a_1 and a_k , where $1 \leq k \leq n$ and k is arbitrary. By anonymity, for all $\pi \in \Pi$ and all $g \in G$,

$$\mathcal{A}(g \cdot \pi) = g \cdot \mathcal{A}(\pi)$$

which immediately yields for all r, j

$$D_{g \cdot \pi, \mathcal{A}}(g(r), j) = D_{\pi, \mathcal{A}}(r, j). \quad (1)$$

Suppose that $1 \leq r \leq n$ and let g be the permutation of $\mathcal{A}$ that simply swaps a_1 and a_r . Then for

each j such that $1 \leq j \leq n$,

$$\begin{aligned}
E_P[D_{\pi, \mathcal{A}}(r, j)] &= \sum_{\pi \in \Pi} P(\pi) D_{\pi, \mathcal{A}}(r, j) && \text{by definition of expectation} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{g \cdot \pi, \mathcal{A}}(g(r), j) && \text{by (1)} \\
&= \sum_{\pi \in \Pi} P(g \cdot \pi) D_{g \cdot \pi, \mathcal{A}}(1, j) && \text{by anonymity of } P \\
&= \sum_{\pi' \in \Pi} P(\pi') D_{\pi', \mathcal{A}}(1, j) && \text{because } g \text{ induces a bijection of } \Pi(A, O) \\
&= E_P[D_{\pi, \mathcal{A}}(1, j)] && \text{by definition of expectation.}
\end{aligned}$$

C Proof of Propostion 4.5

Assuming (i), take h to be the identity to get anonymity and g to be the identity to get neutrality, yielding (ii). Assuming (ii), let π, π' lie in the same orbit so that $\pi' = h \star (g \cdot \pi)$ for some $g \in G, h \in H$. Then $P(\pi') = P(g \cdot \pi)$ by neutrality and this value also equals $P(\pi)$ by anonymity, yielding (iii). Assuming (iii), let $\pi \in \Pi, g \in G, h \in H$. Then $g \cdot (h \star \pi)$ lies in the orbit of π under $G \times H$ and so $P(g \cdot (h \star \pi)) = P(\pi)$. This yields (i).

D Proof of Theorem 4.6

We first give the proof, then an example to clarify the main idea.

Proof. If $g \in G$ then permuting the agents *and the endowment* by g will also permute the assignment given by TTC by g , since the output is entirely determined by the endowment and the preferences. More formally, let $\mathcal{A}_\alpha$ be TTC with endowment α . Then

$$\mathcal{A}_{g \cdot \alpha}(g \cdot \pi) = g \cdot \mathcal{A}_\alpha(\pi).$$

Thus if we write $D_{\pi, \mathcal{A}}$ as $D_{\pi, \alpha}$ when $\mathcal{A}$ is TTC with endowment α , we obtain for all r, j

$$D_{\pi, \alpha}(r, j) = D_{g \cdot \pi, g \cdot \alpha}(g(r), j). \quad (2)$$

Now relabeling the objects back to what they were will restore the original endowment α . Formally, define $\bar{g} \in H$ as the mapping that takes $\alpha(g^{-1}(k))$ to $\alpha(k)$. Then we have $\bar{g} \star g \cdot \alpha = \alpha$. Since relabeling the objects does not change the rank of any agent's allocation, only the label, we obtain

$$D_{g \cdot \pi, g \cdot \alpha}(g(r), j) = D_{\bar{g} \star g \cdot \pi, \alpha}(g(r), j). \quad (3)$$

Thus for each r , if we choose g so that $g(r) = 1$ we obtain

$$\begin{aligned}
E_P[D_{\pi, \alpha}(r, j)] &= \sum_{\pi \in \Pi} P(\pi) D_{\pi, \alpha}(r, j) && \text{by definition of expectation} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{\bar{g} \star (g \cdot \pi), \alpha}(g(r), j) && \text{by (2) and (3)} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{\bar{g} \star (g \cdot \pi), \alpha}(1, j) && \text{by choice of } g \\
&= \sum_{\pi' \in \Pi} P(\pi') D_{\bar{g} \star \pi', \alpha}(1, j) && \text{by anonymity of } P \\
&= \sum_{\pi'' \in \Pi} P(\pi'') D_{\pi'', \alpha}(1, j) && \text{by neutrality of } P \\
&= E_P[D_{\pi, \alpha}(1, j)] && \text{by definition of expectation.}
\end{aligned}$$

□

Example D.1. Consider again the profile π given by

$$\begin{aligned} a_1 : o_1 &\succ o_2 \succ o_3 \\ a_2 : o_1 &\succ o_3 \succ o_2 \\ a_3 : o_2 &\succ o_1 \succ o_3 \end{aligned}$$

and the initial allocation α sending a_i to o_i for each i . Note that with endowment α , the assignment made by TTC on π is $a_1 : o_1, a_2 : o_3, a_3 : o_2$, and the corresponding rank distribution matrix is $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Let g be the 3-cycle $a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow a_1$. The bijections in the above proof yield the following. First, running TTC with endowment $g \cdot \alpha$ on $g \cdot \pi$ we obtain $a_1 : o_2, a_2 : o_1, a_3 : o_3$, and the rank distribution matrix is permuted appropriately: $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Next, $\bar{g}$ is the 3-cycle $o_1 \rightarrow o_2 \rightarrow o_3 \rightarrow o_1$, so that the endowment is transformed back to $a_1 : o_1, a_2 : o_2, a_3 : o_3$ and the profile to

$$\begin{aligned} a_1 : o_3 &\succ o_2 \succ o_1 \\ a_2 : o_2 &\succ o_3 \succ o_1 \\ a_3 : o_2 &\succ o_1 \succ o_3. \end{aligned}$$

The assignment under TTC with endowment α is $a_1 : o_3, a_2 : o_2, a_3 : o_1$ and the rank distribution matrix is still $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Remark D.2. Note that the proof of Theorem 4.6 fails for SD with fixed picking order ρ . The analog of (2) does hold. If we permute ρ and the profile by g , we get the permutation of the original assignment. Letting A_ρ denote SD with picking order ρ , we have $\mathcal{A}_{g \star \rho}(g \cdot \pi) = g \cdot \mathcal{A}_\rho(\pi)$ and so $D_{\pi, \rho}(r, j) = D_{g \cdot \pi, g \star \rho}(g(r), j)$. However the right side of this last equation corresponds to SD with a different order of agents from the original choosing order, and there is no way to restore the original order using only permutations of the objects. That is, there is no analog of (3).

E Proof of Lemma 4.11

We begin the section with the proof, followed by an example illustrating the main concepts.

Proof. For the purposes of this proof, we denote RMM with tiebreak order ρ by $\mathcal{A}_\rho$. Order the objects and agents and make an identification of G and H . That is, choose any bijection $\alpha : \mathcal{A} \rightarrow \mathcal{O}$, and thereby obtain a bijection from G to H given by $g \rightarrow \bar{g} := \alpha \circ g \circ \alpha^{-1}$, where $\circ$ denotes function composition.

Write $D_{\pi, \rho}$ to denote $D_{\pi, \mathcal{A}_\rho}$. Note that as in Remark D.2, with $\rho = \pi(i)$

$$\mathcal{A}_{g \star \rho}(g \cdot \pi) = g \cdot \mathcal{A}_\rho(\pi)$$

so that for all r, j

$$D_{\pi, \rho}(r, j) = D_{g \cdot \pi, g \star \rho}(g(r), j). \quad (4)$$

Also, permuting the objects does not change the rank of the assigned object, only its name. Thus for every $h \in H$ and all r, j

$$D_{\pi, \rho}(g(r), j) = D_{h \star \pi, \rho}(g(r), j). \quad (5)$$

Now for each r , choose g so that $g(r) = 1$ and choose h so that

$$[h \star g \cdot \pi](i) = \bar{g} \star \pi(i).$$

Note that this is possible because $[g \cdot \pi](i)$ and $\bar{g} \star \pi(i)$ are both strict linear orders over the objects and H acts transitively (there is a single orbit).

We obtain

$$\begin{aligned}
E_P[D_{\pi, \alpha^{-1}\pi(i)}(r, j)] &= \sum_{\pi \in \Pi} P(\pi) D_{\pi, \alpha^{-1}\pi(i)}(r, j) && \text{by definition of expectation} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{g \cdot \pi, g \star \alpha^{-1}\pi(i)}(g(r), j) && \text{by (4)} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{g \cdot \pi, \alpha^{-1}\bar{g} \star \pi(i)}(g(r), j) && \text{by definition of } \bar{g} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{h \star g \cdot \pi, \alpha^{-1}\bar{g} \star \pi(i)}(g(r), j) && \text{by (5)} \\
&= \sum_{\pi \in \Pi} P(\pi) D_{h \star g \cdot \pi, \alpha^{-1}[h \star g \cdot \pi](i)}(g(r), j) && \text{by choice of } h \\
&= \sum_{\pi \in \Pi} P(\pi) D_{h \star g \cdot \pi, \alpha^{-1}[h \star g \cdot \pi](i)}(1, j) && \text{by choice of } g \\
&= \sum_{\pi' \in \Pi} P(\pi') D_{h \star \pi', \alpha^{-1}h \star \pi'(i)}(1, j) && \text{by anonymity of } P \\
&= \sum_{\pi'' \in \Pi} P(\pi'') D_{\pi'', \alpha^{-1}\pi''(i)}(1, j) && \text{by neutrality of } P \\
&= E_P[D_{\pi, \alpha^{-1}\pi(i)}(1, j)] && \text{by definition of expectation.}
\end{aligned}$$

□

Example E.1. The bijections in the last proof can be demonstrated using Example 4.10, with $\alpha(a_i) = o_i$ for each i . Let π be the profile in that example, let the tiebreak order $\rho = a_3 > a_1 > a_2$ determined by the first agent (in other words $i = 1$), and let g be the cycle $a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow a_1$. Then $g \cdot \pi$ is the profile

$$\begin{aligned}
a_1 : o_2 &\succ o_3 \succ o_1 \\
a_2 : o_3 &\succ o_1 \succ o_2 \\
a_3 : o_3 &\succ o_1 \succ o_2
\end{aligned}$$

and the two rank maximal matchings are $a_1 : o_2, a_2 : o_1, a_3 : o_3$ and $a_1 : o_2, a_2 : o_3, a_3 : o_1$. With the picking order $g \star \rho = a_1 > a_2 > a_3$ the assignment is $a_1 : o_2, a_2 : o_3, a_3 : o_1$, so that, for example, agent $a_2 = g(a_1)$ receives her first choice, agent $a_3 = g(a_2)$ receives her second choice, and agent $a_1 = g(a_3)$ receives her first choice.

We must choose h to satisfy $h \star (o_2, o_3, o_1) = (o_1, o_2, o_3)$ so that h is the 3-cycle $o_1 \rightarrow o_3 \rightarrow o_2 \rightarrow o_1$. The profile $h \star g \cdot \pi$ is

$$\begin{aligned}
a_1 : o_1 &\succ o_2 \succ o_3 \\
a_2 : o_2 &\succ o_3 \succ o_1 \\
a_3 : o_2 &\succ o_3 \succ o_1
\end{aligned}$$

so with the picking order $a_1 > a_2 > a_3$ the mechanism outputs the assignment $a_1 : o_1, a_2 : o_2, a_3 : o_3$. Again, as they must, agents a_1 and a_2 receive their first choice and a_3 its second.

F Proof of Theorem 4.12

We showed order symmetry in Lemma 4.11, and ex ante anonymity follows from Appendix H.2. For ordinal efficiency, we show that for any preference profile π , the fractional assignment $\text{RMM-RA}(\pi)$ is rank-maximal in the space of fractional assignments. Since a rank-maximal matching is welfare maximizing when agents have lexicographic utilities, any rank-maximal matching satisfies rank efficiency (Featherstone 2020), a strengthening of ordinal efficiency.

We proceed to show rank maximality. To see this, suppose to the contrary that there exists a fractional assignment α and associated rank distribution matrix D_α for which there exists an index j for which

$$\sum_{i=1}^n D_{\pi, \text{RMM-RA}}(i, j') = \sum_{i=1}^n D_\alpha(i, j')$$

for all $j' < j$, and

$$\sum_{i=1}^n D_{\pi, \text{RMM-RA}}(i, j) < \sum_{i=1}^n D_{\alpha}(i, j).$$

That is, the expected number of agents who receive their j' th most-preferred object is equal for both $\text{RMM-RA}(\pi)$ and α for all $j' < j$, but the expected number of agents who receive their j th most-preferred object is higher for α than for $\text{RMM-RA}(\pi)$.

Further, for every rank index j' we know that for every tiebreak order,

$$\sum_{i=1}^n D_{\pi, \text{RMM}}(i, j') = \sum_{i=1}^n D_{\pi, \text{RMM-RA}}(i, j'),$$

because every output of RMM is a rank-maximal matching, which uniquely determines the distribution of ranks received by the agents.

Consider any Birkhoff-von Neumann decomposition of α into discrete assignments $\alpha_1, \dots, \alpha_k$ such that $\alpha = \sum_{\ell=1}^k q_{\ell} \alpha_{\ell}$ and corresponding rank distribution matrices $D_{\alpha_1}, \dots, D_{\alpha_k}$. Let $j' < j$. Then $\sum_{\ell=1}^k q_{\ell} \sum_{i=1}^n D_{\alpha_{\ell}}(i, j') = \sum_{i=1}^n D_{\alpha}(i, j')$. It must be the case that

$$\sum_{i=1}^n D_{\alpha_{\ell}}(i, j') = \sum_{i=1}^n D_{\alpha}(i, j') \quad (6)$$

for all $\ell \in \{1, \dots, k\}$. If not, consider the smallest such j' for which Equation 6 is violated. Then there exists an ℓ' for which

$$\sum_{i=1}^n D_{\alpha_{\ell'}}(i, j') > \sum_{i=1}^n D_{\alpha}(i, j') = \sum_{i=1}^n D_{\pi, \text{RMM-RA}}(i, j') = \sum_{i=1}^n D_{\pi, \text{RMM}}(i, j'),$$

violating rank-maximality of $\text{RMM}(\pi)$, since $\alpha_{\ell'}$ and $\text{RMM}(\pi)$ allocate the same number of agents to each of their top $j' - 1$ objects, but α allocates more agents to their j' th most-preferred object than $\text{RMM}(\pi)$ does.

We therefore arrive at rank index j , where we have

$$\sum_{\ell=1}^k q_{\ell} \sum_{i=1}^n D_{\alpha_{\ell}}(i, j) = \sum_{i=1}^n D_{\alpha}(i, j) > \sum_{i=1}^n D_{\pi, \text{RMM-RA}}(i, j) = \sum_{i=1}^n D_{\pi, \text{RMM}}(i, j),$$

which is a contradiction of rank-maximality of $\text{RMM}(\pi)$ by the same argument as applied immediately above.

G Decomposing assignment mechanisms

By the Birkhoff-von Neumann theorem, every assignment can be expressed as the convex combination of discrete assignments. Let us define a matrix with $(n!)^n$ rows and $n!$ columns, where each row corresponds to a profile and each column corresponds to a discrete assignment. Note that any assignment mechanism can be expressed as a row stochastic matrix, which corresponds to a mapping from profiles to distributions over discrete assignments. Further, any row stochastic matrix can be expressed as a convex combination of binary row stochastic matrices in which every entry is either 0 or 1. Note that any binary row stochastic matrix corresponds to a discrete assignment mechanism, since it maps each profile to a single assignment. Therefore, any assignment mechanism can be expressed as a convex combination of discrete assignment mechanisms.

H Group actions

We use group-theoretic concepts in several proofs, and first recall some standard abstract facts. A *left action* of a group G on a set X is a mapping $G \times X \rightarrow X$, often denoted $(g, x) \mapsto g \cdot x$, such that

- $1 \cdot x = x$ for all $x \in X$;
- $h \cdot (g \cdot x) = (hg) \cdot x$ for all $h, g \in G$ and all $x \in X$.

Alternatively, it is a homomorphism from G into the group $\text{Sym}(X)$ of all permutations of X . We use left actions because the more common convention for function composition is to compose from right to left.

We can build up actions on larger sets in standard ways. Suppose that G acts on X and Y is a set. Then there is an action of G on the set Z of functions from Y to X , given by

$$[g \star f](y) = g \cdot f(y).$$

There is also an action of G on the set of functions from X to Y , given by

$$[g \cdot f](x) = f(g^{-1} \cdot x).$$

The g^{-1} is so that we indeed get a left action:

$$[h \cdot [g \cdot f]](x) = (g \cdot f)(h^{-1}(x)) = f(g^{-1}(h^{-1}(x))) = f((hg)^{-1}(x)) = (hg) \cdot f(x).$$

Note that we can also write

$$[g \cdot f](g \cdot x) = f(x).$$

Specializing to our situation, we apply these constructions in several places. First, a linear order of objects is simply a function from $Y := \{1, \dots, n\}$ to $\mathcal{O}$. Thus the group $H := \text{Sym}(\mathcal{O})$ acts via

$$[h \star \sigma](i) = h(\sigma(i)).$$

Similarly, an ordering of agents is a function from Y to $\mathcal{A}$, and we have

$$[g \star \rho](i) = g(\rho(i)).$$

An assignment is a mapping from $\mathcal{A}$ to $\mathcal{O}$, so we can use both constructions. We have

$$[g \cdot \alpha](a) = \alpha(g^{-1} \cdot a)$$

so that $g(a)$ is assigned the object that a had under α , and

$$[h \star \alpha](a) = h(\alpha(a)).$$

Going to the next level, we see that a profile is a mapping from $\mathcal{A}$ to $L(\mathcal{O})$. Thus we have

$$[g \cdot \pi](a) = \pi(g^{-1}(a))$$

which says that in the profile $g \cdot \pi$, $g(a)$ has the preference order that a had in π ; that is, the preference orders have been permuted consistently. Similarly

$$[h \star \pi](a) = h(\pi(a))$$

which consistently relabels the objects in each agent's preference order.

H.1 Anonymity and neutrality

An anonymous mechanism is one for which

$$\mathcal{A}(g \cdot \pi) = g \cdot \mathcal{A}(\pi)$$

for all $g \in G$, and a neutral mechanism is one for which

$$\mathcal{A}(h \star \pi) = h \star \mathcal{A}(\pi).$$

In each case, a consistent permutation of preferences yields the analogous permutation of the assignment.

Anonymity is a strong condition that cannot be satisfied by discrete assignment mechanisms. More commonly we have a parametrized family of mechanisms, and we can derive some useful formulae. For example, SD and RMM have a parameter ρ , the choosing order of agents, while TTC has a parameter α , the initial assignment. Each of these parameters depends on the agents, and agents play different roles, indicating that anonymity is not satisfied.

If the mechanism's output depends only the preferences and the agent roles, and not on the agent's identities, then a consistent relabeling of agents will yield a consistent relabeling of the output assignment. More precisely,

$$SD_{g \star \rho}(g \cdot \pi) = g \cdot SD_{\rho}(\pi) \quad (7)$$

$$RMM_{g \star \rho}(g \cdot \pi) = g \cdot SD_{\rho}(\pi) \quad (8)$$

$$TTC_{g \cdot \alpha}(g \cdot \pi) = g \cdot TTC_{\alpha}(\pi). \quad (9)$$

In each case, on the left agent $g(a)$ plays the same role and has the same preferences as a does on the right.

H.2 Symmetrization

We now show that given any discrete assignment mechanism $\mathcal{A}$, we can define a corresponding randomized assignment mechanism $\overline{\mathcal{A}}$ whose associated fractional assignment mechanism $\overline{\mathcal{A}}$ is anonymous. We note that Bade (2020) defines the same construction, denoting it $\Delta\mathcal{A}$ and naming it the *symmetrization* of $\mathcal{A}$.

Proposition H.1. *Let $\mathcal{A}$ be a discrete assignment mechanism. For each $g \in G$, define*

$$\mathcal{A}^g(\pi) := g \cdot \mathcal{A}(g^{-1} \cdot \pi).$$

Then $\mathcal{A}^g$ is a discrete assignment mechanism. Furthermore we can define a randomized mechanism $\overline{\mathcal{A}}$ whose associated fractional mechanism $\overline{\mathcal{A}}$ given by

$$\overline{\mathcal{A}} := \frac{1}{|G|} \sum_{g \in G} \mathcal{A}^g$$

is anonymous.

Proof. Each $\mathcal{A}^g$ clearly outputs a discrete assignment and the weights $1/|G|$ sum to 1, so $\overline{\mathcal{A}}$ and $\overline{\mathcal{A}}$ are well defined and of the correct type. Now let $h \in G$. Letting $k = h^{-1}g$ below,

$$\begin{aligned} \overline{\mathcal{A}}(h \cdot \pi) &= \frac{1}{|G|} \sum_{g \in G} \mathcal{A}^g(h \cdot \pi) = \frac{1}{|G|} \sum_{g \in G} g \cdot \mathcal{A}(g^{-1} \cdot (h \cdot \pi)) \\ &= \frac{1}{|G|} \sum_{g \in G} g \cdot \mathcal{A}((g^{-1}h) \cdot \pi) = \frac{1}{|G|} \sum_{k \in G} (hk) \cdot \mathcal{A}(k^{-1} \cdot \pi) \\ &= h \cdot \left(\frac{1}{|G|} \sum_{k \in G} k \cdot \mathcal{A}(k^{-1} \cdot \pi) \right) = h \cdot \overline{\mathcal{A}}(\pi). \end{aligned}$$

□

Note that $\mathcal{A}$ is anonymous if and only if $\mathcal{A}^g = \mathcal{A}$ for each $g \in G$. Also $\mathcal{A}^g$ is equivalently defined by the identity

$$g^{-1} \cdot \mathcal{A}^g(\pi) = \mathcal{A}(g^{-1} \cdot \pi).$$

Example H.2. *Suppose that $\mathcal{A}$ is SD with a fixed picking order $\rho = (a_1, \dots, a_n)$. Then $\mathcal{A}^g$ is SD with the picking order $(g(a_1), \dots, g(a_n))$, and $\overline{\mathcal{A}}$ is **RSD**. This is because transforming to $g^{-1} \cdot \pi$ has the effect of changing the preference of a_1 to that of $g(a_1)$, etc. Thus a_1 now gets what $g(a_1)$ ranked first, etc, and then applying g ensures that $g(a_1)$ gets its first choice, etc. More precisely, writing $g \star \rho = \mu$ we have*

$$SD_{g^{-1} \star \mu}(g^{-1} \cdot \pi) = g^{-1} \cdot SD_{\mu}(\pi)$$

so that $\mathcal{A}^g = SD_{\mu}$.

*The same argument works for RMM with a fixed tiebreaking order. Here $\mathcal{A}^g$ is RMM with the tiebreaking agent transformed by g , and $\overline{\mathcal{A}}$ is **RMM-RA**.*

*Similarly, if $\mathcal{A}$ is TTC with a fixed endowment α , $\mathcal{A}^g$ is TTC with endowment $g \cdot \alpha$, and $\overline{\mathcal{A}}$ is **TTC-RE**. The analogous argument works: writing $g \cdot \alpha = \beta$ we have*

$$TTC_{g^{-1} \cdot \beta}(g^{-1} \cdot \pi) = g^{-1} \cdot TTC_{\beta}(\pi)$$

so that $\mathcal{A}^g = TTC_{\beta}$.

H.3 Coupling of actions

Consider again the actions $\cdot$ and $\star$ on assignments and profiles. The actions commute with each other, because permuting agents and permuting objects are completely independent:

$$\begin{aligned} h \star (g \cdot \pi) &= g \cdot (h \star \pi) \\ h \star (g \cdot \alpha) &= g \cdot (h \star \alpha). \end{aligned}$$

In our situation where $|\mathcal{A}| = |\mathcal{O}|$, the groups $\text{Sym}(\mathcal{A})$ and $\text{Sym}(\mathcal{O})$ are different as groups, but they are each isomorphic to the group S_n of all permutations of $\{1, 2, \dots, n\}$. Every choice of a linear order of agents (respectively objects) gives an isomorphism. We call this “identifying G with S_n ”. If we do this for both G and H then we can identify G with H .

Another way of writing this is to fix a bijection $\alpha : \mathcal{A} \rightarrow \mathcal{O}$ (which is simply an assignment). Then there is an isomorphism from G to H given by

$$\bar{g} = \alpha \circ g \circ \alpha^{-1} \in H$$

with inverse

$$h \mapsto \alpha^{-1} \circ h \circ \alpha \in G.$$

It follows immediately that

$$[\bar{g} \star [g \cdot \alpha]](a) = \bar{g}([g \cdot \alpha](a)) = \bar{g}(\alpha(g^{-1}(a))) = \alpha(a). \quad (10)$$

Similarly

$$[\bar{g} \star [g \cdot \pi]](g(a)) = \bar{g} \star \pi(a). \quad (11)$$

Thus transforming the profile and agent by g , and then transforming the objects by $\bar{g}$, is the same as just transforming the objects.

We can relate the two actions as follows. For a given g and π , there is in general no way to express the permutation of agents $g \cdot \pi$ simply by permuting the objects via $h \star \pi$. However we can at least do this for a single agent. For each g, π, a we can find h so that

$$[g \cdot \pi](a) = h \star \pi(a).$$

This is clear because the left side is simply a linear order (permutation) on objects, and we can obtain any other linear order by an appropriate permutation of objects.

Example H.3. Consider again the profile π used in Examples 3.2, 3.4 and ?? with endowment α for which $\alpha(a_i) = o_i$ for all i . Let g be the 3-cycle $a_1 \rightarrow a_2 \rightarrow a_3 \rightarrow a_1$ as before. The relabeling of agents by g takes π to the profile $g \cdot \pi$ given by

$$\begin{aligned} a_1 : o_2 \succ o_1 \succ o_3 \\ a_2 : o_1 \succ o_2 \succ o_3 \\ a_3 : o_1 \succ o_3 \succ o_2 \end{aligned}$$

and the endowment to $g \cdot \alpha$ given by $a_1 : o_3, a_2 : o_1, a_3 : o_2$. The relabeling of objects by $\bar{g}$ moves the endowment back to α and the profile to $\bar{g} \star g \cdot \pi$ given by

$$\begin{aligned} a_1 : o_3 \succ o_2 \succ o_1 \\ a_2 : o_2 \succ o_3 \succ o_1 \\ a_3 : o_2 \succ o_1 \succ o_3. \end{aligned}$$

I Example of Anonymity

In the profile $g \cdot \pi$, agent $a_{g(i)}$ has the preference order that agent a_i had in π . Anonymity of $\mathcal{A}$ says that on this input, $\mathcal{A}$ gives agent $a_{g(i)}$ what agent a_i obtained under π .

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

$$\begin{aligned} a_1 : o_1 \succ o_2 \succ o_3 \\ a_2 : o_1 \succ o_3 \succ o_2 \\ a_3 : o_2 \succ o_1 \succ o_3 \end{aligned}$$

and suppose that the assignment α made by some algorithm is

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

Taking g to be the 3-cycle mapping $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$, we obtain $g \cdot \pi$

$$\begin{aligned} a_1 : o_2 \succ o_1 \succ o_3 \\ a_2 : o_1 \succ o_2 \succ o_3 \\ a_3 : o_1 \succ o_3 \succ o_2 \end{aligned}$$

and $g \cdot \alpha$

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

J Order Symmetry for Non-Anonymous Measures

There exists a fractional mechanism that is order symmetric with respect to all probability measures. Equivalently, there exists a randomized assignment mechanism that is ex ante order symmetric with respect to all measures.

Example J.1. Consider the assignment mechanism $\mathcal{A}$ that assigns each agent a $\frac{1}{n}$ fraction of each object, independent of preferences. This is order-symmetric for every P , because $D_{\pi, \mathcal{A}}(r, j) = \frac{1}{n}$ for all profiles π and all r, j .

Clearly, this is a degenerate mechanism that fails to take advantage of differences in the agents' preferences. It is therefore natural to wonder whether order symmetry and ordinal efficiency can be achieved simultaneously. Characterizing the measures for which this combination can be achieved remains an interesting open question. However, we show that there exist measures for which it is impossible.

Example J.2. Consider the probability measure P with all its weight on the following profile π :

$$\begin{aligned} a_1, \dots, a_{n-1} : o_1 \succ o_2 \succ \dots \succ o_n \\ a_n : o_2 \succ o_1 \succ \dots \succ o_n \end{aligned}$$

Order symmetry imposes that $D_{\pi, \mathcal{A}}(r, j)$ is independent of r . That is, all agents have the same probability of receiving their j th most-preferred object, for all j . In this instance it must be the case that $D_{\pi, \mathcal{A}}(r, 1) = 1/(n-1)$ for any order-symmetric $\mathcal{A}$; that is, each agent receives their top choice with probability $1/(n-1)$. Clearly $D_{\pi, \mathcal{A}}(r, 1) > 1/(n-1)$ is impossible, as then object o_1 would be over-allocated. Further, if $D_{\pi, \mathcal{A}}(r, 1) < 1/(n-1)$ then agent a_n receives object o_1 with some positive probability, and therefore, by order symmetry, agents $a_1, \dots, a_{n-1}$ receive object o_2 with some positive probability, violating ordinal efficiency.

However, this leaves some positive probability for agents $a_1, \dots, a_{n-1}$ to be assigned object o_2 , but zero probability that agent a_n is assigned object o_1 , violating order symmetry for $j = 2$.

K Order Bias of Common Algorithms for a Special Case

Let $1 \leq k \leq n$ and let P_k be any probability measure whose support consists entirely of profiles for which agents $1, \dots, k$ have the same preference order (without loss of generality $1 \succ 2 \succ \dots \succ n$), but the last $n - k$ agents have different first choices taken from the $n - k$ bottom choices of the first k agents. For example, when $k = 1$ or $k = n$ we can obtain the measures discussed in Example 4.1. On each profile in the support of P_k , SD or either variant of Boston, if run with the choosing order $1, 2, \dots, n$, will assign agent i its i th choice for $1 \leq i \leq k$ and the other agents their first choice. The same is true of TTC if run with initial endowment where agent i has object i . Thus the order bias of each of these algorithms with respect to P_k and scoring rule s is $1 - s_k$.

L Omitted material from Section A.2

L.1 Theoretical results relating Naive and Adaptive Boston

In our simulation results, we see that Naive Boston (NB) consistently outperforms Adaptive Boston (AB) in terms of order bias. To what extent this is generally true remains an open question, but we are able to show that NB achieves lower order bias than AC with respect to IC for any scoring rule that has non-zero entries in only the first two positions (with plurality and 2-approval being natural examples).

Lemma L.1. *Consider NB and AB both with picking order equal to $\{a_1, \dots, a_n\}$.*

1. *Agent a_n receives their top-choice object under NB on a profile π if and only if they receive their top-choice object under AB.*
2. *If agent a_n receives their second-choice object under AB on a profile π then they also receive their second-choice object under NB.*
3. *There exists a profile π' in which a_n receives their second-choice object under NB but a less preferred object under AB.*

Proof. For the first part, note that under both AB and NB agent a_n gets their top choice if and only if no other agent has the same top choice (since a_n is last in the tiebreak order, all other agents have higher priority).

For the second part, suppose that a_n receives their second-choice object (call it o_j) under AB on profile π . Then it must be the case that no other agent has o_j as their top-choice object (or else that agent would receive o_j in the first round). Further, it cannot be the case that any other agent submits a bid for o_j in the second round (since such an agent would be awarded o_j ahead of a_n). In particular, it cannot be the case that any other agent who has o_j as their second choice is not awarded their first choice object. But this implies that no agent other than a_n submits a bid for o_j in the second round of the NB mechanism. Thus, agent a_n receives o_j under NB on profile π .

Finally, consider profile π' where agents a_1, a_2, a_3, a_4 have respective preferences over o_1, o_2, o_3, o_4 as follows.

$$\begin{aligned} a_1 : o_1 &\succ o_2 \succ o_3 \succ o_4 \\ a_2 : o_1 &\succ o_2 \succ o_3 \succ o_4 \\ a_3 : o_2 &\succ o_1 \succ o_3 \succ o_4 \\ a_4 : o_1 &\succ o_3 \succ o_2 \succ o_4 \end{aligned}$$

Under NB, in the first round a_1 receives o_1 and a_3 receives o_2 . In the second round, a_2 bids on o_2 and a_4 bids on (and receives) o_3 . However, under AB, a_2 bids on (and receives) o_3 in the second round since their second choice o_2 has already been allocated in the first round, with a_4 receiving o_4 instead. $\square$

Lemma L.1 readily yields the following result.

Theorem L.2. *Let s be a positional scoring rule with $s_3 = 0$. Then the order bias of NB is less than that of AB with respect to scoring rule s and the IC measure on profiles. That is, $\beta_n(NB, ; s; IC) \leq \beta_n(AB, ; s; IC)$. If $s_2 > 0$ then the inequality is strict.*

Proof. Under both NB and AB, the first agent in the tiebreak order always gets their most preferred object. Thus the order bias is dictated by the expected utility of the agent with lowest expected utility. Under IC it is clear that the last agent in the tiebreak order receives the lowest expected utility. By Lemma L.1, the last agent gets weakly higher utility for every profile under NB than they do under AB, which establishes the weak inequality. The fact that the inequality is strict whenever $s_2 > 0$ follows immediately from the third part of Lemma L.1. $\square$

Note that the only property of the IC measure that is used in the proof of Theorem L.2 is that the last agent in the tiebreak order receives the lowest expected utility under both NB and AB. In fact, this property continues to hold for any anonymous measure (Pritchard and Wilson 2022), and therefore so does Theorem L.2.

It remains open to what extent Theorem L.2 further generalizes to other probability measures and to more general scoring rules. However, it is not the case that NB has lower order bias than AB in all situations. In fact, in the following example, the measure on profiles is anonymous and AB has lower

order bias than NB for all scoring rules. Therefore, any explanation for the lower order bias of NB under impartial culture and Mallows preferences for scoring rules such as Borda and antiplurality must exploit some additional property of those measures.

Example L.3. Suppose that $n = 5$, with a probability measure on profiles defined as follows. There are two possible preference orders: $o_1 \succ o_2 \succ o_3 \succ o_4 \succ o_5$ and $o_2 \succ o_5 \succ o_3 \succ o_1 \succ o_4$, with associated probabilities p and $1 - p$. Each agent randomly chooses a preference order, independently of other agents. So, there are 32 possible profiles. For 28 of the profiles, NB and AB agree on which item should be assigned to the last agent. For the other 4 profiles, AB assigns the last agent in the tiebreak order their third choice, while NB assigns the last agent their fourth choice.

L.2 Empirical results for other utility functions

Borda utilities give weight to all ranks in a linearly decreasing way. In Figures 4–6 we present analogous results to those for Borda utilities in the main text, for two very different utility functions. The first is induced by the 2-approval scoring rule $(1, 1, 0, 0, \dots, 0)$. Thus we measure the fraction of profiles for which each agent gets one of their top two choices. By the results of the previous subsection, we know that NB achieves lower order bias than AB for this scoring rule and any anonymous measure, and we accordingly observe this in the results that follow. The second is induced by the antiplurality scoring rule $(1, 1, \dots, 1, 0)$, and measures the fraction of profiles for which each agent gets anything better than their last choice.

For large n and Impartial Culture, the Boston mechanisms usually manage to give most agents one of their top few choices, leading to the declining order bias as measured by Borda (Fig. 1) and antiplurality (Fig. 4, right). Even the last agent has a non-zero limiting probability (≈ 0.5) of obtaining one of the first two preferences (Fig. 4, left). Similar results are obtained for the Mallows model with large dispersion ($\phi \approx 1$).

For the Mallows model with small dispersion ($\phi = 0.6$ is illustrated in Fig. 5), the situation is closer to that in which all agents have the same preference order. For large n , the last agent is unlikely to receive one of the first two preferences, and may even have non-zero limiting probability of receiving the very last preference.

Figure 6 examines the performance of the algorithms in terms of order bias while varying the Mallows parameter from 0 to 1. For both 2-approval and antiplurality scoring rules, NB is a strong choice for almost all values of ϕ .




Figure 4: (left) 2-approval order bias, Impartial Culture, mean of 1000 simulations, (right) antiplurality order bias, Impartial Culture, mean of 1000 simulations.

L.3 Empirical results for variants of the Boston algorithms

The first agent clearly has a consistent advantage over others, and particularly over the last agent, in both Naive and Adaptive Boston. An obvious way to try to reduce order bias is to reverse the choosing order of the agents at each round. Thus the last agent will do much better on average than in the original algorithms. In fact the identity of the worst-off agent is no longer clear and depends on the utility function. Figures 7–11 are analogous to those in Sections A.2 and L.2. They show that the reversed




Figure 5: (left) 2-approval order bias, Mallovs ($\phi = 0.6$), mean of 1000 simulations, (right) antiplurality order bias Mallovs ($\phi = 0.6$), mean of 1000 simulations.




Figure 6: (left) 2-approval order bias, $n = 32$, mean of 1000 simulations, (right) antiplurality order bias, $n = 32$, mean of 1000 simulations.

forms (shown in the legends as Z for the naive case and X for the adaptive case) indeed have lower order bias than the originals. Throughout, the reversing form of Naive Boston has the best performance.

L.4 Empirical results on welfare

In Figures 12 and 13 we present the utilitarian and Nash welfare (with respect to Borda utilities), again averaged over 1000 independent samples. In Figures 14 and 15 we repeat the experiment of utilitarian welfare with respect to the other utility functions.

For large n and Impartial Culture, all the mechanisms usually manage to give most agents one of their top few choices. The very last agent in a Serial Dictatorship mechanism is an exception – but as only one agent among many, he has little effect on overall welfare. The Borda or anti-plurality welfare thus approaches 1, and even the 2-approval welfare approaches a respectably high limit (depending on the mechanism; Figures 12, 14).

For the Mallovs model with small dispersion ($\phi = 0.6$ is illustrated in Figures 13, 15), the average agent fares less well: the 2-approval and even Borda welfare expectations are well below 1. Note that in Figure 13, in contrast to Figure 12, both welfare measures decrease with n , and that the results are highly dependent on the utility function. Overall, the relative performance of the mechanisms is similar to Figure 12.

The Boston mechanisms exhibit a small advantage in terms of utilitarian welfare but a disadvantage in terms of Nash welfare, suggesting that the extra additive welfare they achieve relative to SD and TTC is distributed unevenly within each instance.

Finally, note that with respect to Nash and utilitarian welfare, TTC and SD give almost identical results, as we expect.




Figure 7: (left) Borda order bias, Impartial Culture, mean of 1000 simulations, (right) distribution of rank received for last agent, $n = 16$, Impartial Culture, mean of 1000 simulations.




Figure 8: (left) Borda order bias, Mallows(0.6), mean of 1000 simulations, (right) distribution of rank received for last agent, $n = 16$, Mallows (0.6), mean of 1000 simulations.




Figure 9: (left) 2-approval order bias, Impartial Culture, mean of 1000 simulations, (right) antiplurality order bias, Impartial Culture, mean of 1000 simulations.




Figure 10: (left) 2-approval order bias, Mallovs ($\phi = 0.6$), mean of 1000 simulations, (right) antiplurality order bias Mallovs ($\phi = 0.6$), mean of 1000 simulations.




Figure 11: (left) 2-approval order bias, $n = 32$, mean of 1000 simulations, (right) antiplurality order bias, $n = 32$, mean of 1000 simulations.




Figure 12: (left) Borda utilitarian welfare, Impartial Culture, mean of 1000 simulations, (right) Borda Nash welfare, Impartial Culture, mean of 1000 simulations.




Figure 13: (left) Borda utilitarian welfare, Mallows ($\phi = 0.6$), mean of 1000 simulations, (right) Borda Nash welfare, Mallows ($\phi = 0.6$), mean of 1000 simulations.




Figure 14: (left) 2-approval utilitarian welfare, Impartial Culture, mean of 1000 simulations, (right) antiplurality utilitarian welfare, Impartial Culture, mean of 1000 simulations.




Figure 15: (left) 2-approval utilitarian welfare, Mallows (0.6), mean of 1000 simulations, (right) antiplurality utilitarian welfare, Mallows(0.6), mean of 1000 simulations.