Order Symmetry: A New Fairness Criterion for Assignment Mechanisms

Rupert Freeman University of Virginia Charlottesville, United States freemanr@darden.virginia.edu Geoffrey Pritchard University of Auckland Auckland, New Zealand g.pritchard@auckland.ac.nz

Mark C. Wilson University of Massachusetts Amherst, United States markwilson@umass.edu

ABSTRACT

We introduce a new criterion, order symmetry, for assignment mechanisms that match n objects to n agents having ordinal preferences over the objects. An assignment mechanism is order-symmetric with respect to a given probability measure over preference profiles if every agent has equal probability of receiving their favorite object, equal probability of receiving their second favorite, and so on. Crucially, and unlike other fairness notions such as anonymity or envy-freeness, order symmetry can be satisfied by discrete assignment mechanisms when associated with a sufficiently symmetric probability measure. It can also be interpreted as a criterion of procedural fairness or fairness under uncertainty. 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 efficiency by the Top Trading Cycles mechanism, but not by Serial Dictatorship. We also use the lens of order symmetry to improve the fairness of existing mechanisms with no loss in social welfare, focusing on the widely used family of Boston mechanisms. In addition to theoretical results, we present simulations using data from the Mallows distribution over its full range of parameters, which show an improvement in fairness even on probability measures for which full order symmetry is impossible.

KEYWORDS

one-sided matching; house allocation; fairness

ACM Reference Format:

Rupert Freeman, Geoffrey Pritchard, and Mark C. Wilson. 2025. Order Symmetry: A New Fairness Criterion for Assignment Mechanisms. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

Suppose that two agents Alice and Bob must each receive a fruit, and a banana and apple are available. Each agent has a strict preference order for the fruits. If the agents have different top choices, then each can get their top choice. However if their top choices coincide, only one agent can get it, while the other must get their second choice. If we do not know in advance the agents' preferences over the fruits, we need a mechanism to allocate the fruits.

This work is licensed under a Creative Commons Attribution International 4.0 License. Two common allocation mechanisms are Top Trading Cycles (TTC) and Serial Dictatorship (SD). In the two-agent case as above, SD selects one agent to choose first, leaving the other with the remaining item. If Alice chooses first, she gets her top choice in all four possible preference profiles, while Bob does so only in the two profiles where their preferences differ. Under TTC, items are initially assigned arbitrarily, and agents swap if both prefer the other's item. A simple calculation shows that each agent receives their first choice in three profiles—for example, Bob only fails when preferences are identical (we call this a *unanimous* profile) and Alice is initially assigned the common top choice. Table 1 details these outcomes. Intuitively, TTC produces a fairer collection of allocations than SD.¹

This toy example encapsulates the key aspects of the general case we analyze in this paper. First, the allocation mechanism must be chosen without knowing the agents' preferences. Second, in the absence of such knowledge, all possible preference profiles are considered. Third, the goal is to equalize each agent's chances of achieving good outcomes, avoiding any *a priori* advantage.

The Alice and Bob scenario above 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 [10, 20, 44], operations research [6, 17, 41], and computer science [2, 22, 29, 40]. The model was introduced by Hylland and Zeckhauser [25], who also mention assigning legislators to committees.

In the discrete deterministic model of house allocation described above, some outcome asymmetry is inevitable. With unanimous preferences, one agent gets the best object while another gets the worst. This inherently violates several fairness notions, including anonymity (i.e., symmetric treatment of agents), equal-treatmentof-equals (i.e., agents with identical preferences receive the same allocation), and envy-freeness (i.e., no agent should prefer another's allocation to her own) [22]. However, the Alice and Bob example suggests a new fairness criterion based on average-case analysis. Given a probability measure P on profiles, a mechanism is fair (with respect to P) if all agents do equally well *in expectation with respect to* P. Note that this property makes sense for deterministic mechanisms, and we are discussing something conceptually quite distinct from the randomization in a randomized mechanism.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

¹While this argument implicitly assumes all profiles are equally likely, the conclusion holds for any full-support distribution over the four profiles.

Alice prefs	Bob prefs	SD outcome	SD ranks	TTC outcome	TTC ranks
a > b	a > b	A:a,B:b	A: 1, B: 2	A:a,B:b	A: 1, B: 2
a > b	b > a	A:a,B:b	A: 1, B: 1	A:a,B:b	A: 1, B: 1
b > a	a > b	A:b,B:a	A: 1, B: 1	A:b,B:a	A: 1, B: 1
b > a	b > a	A:b,B:a	A: 1, B: 2	A:a,B:b	A: 2, B: 1

Table 1: Alice, Bob and their fruits. We assume the initial endowment for TTC is A : a, B : b.

This idea clarifies why serial dictatorship and similar mechanisms that rely on a fixed choosing or tiebreak order are unfair some agents consistently fare better across all profiles. Since agents provide only ordinal preferences, we evaluate fairness based on the rank of their assigned object. A mechanism is order-symmetric with respect to P if all agents have equal probabilities of receiving their top object, their second object, and so on, reflecting the intuition that agents are not systematically ordered by the mechanism. If these probabilities differ, the discrepancy serves as a measure of unfairness, which we call the order bias of the mechanism.

In independent and simultaneous work, Long and Velez [32] (based on the results of Long [31]) 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 suggest any new mechanisms. Some other work also uses average-case analysis. Pycia and Ünver [39] 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; we do not consider such preferences in our work. Harless and Manjunath [24] 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 (with expectation taken over uncertainty in the utility functions) more evenly than SD. Other work examining average-case analysis in the one-sided matching setting [18, 21, 23] has focused on social welfare approximation rather than symmetry in agent outcomes. Beyond house allocation, Ozkes and Sanver [36] study the general social choice setting in which voters submit complete preferences over a set of alternatives, and define consequential neutrality, a relaxation of neutrality (i.e., symmetric treatment of alternatives) that requires each alternative to win under the same number of profiles. Consequential neutrality is reminiscent of order symmetry. However, we are more interested in relaxing anonymity because we care about fairness to agents and because, unlike anonymity, neutrality can be satisfeid exactly in house allocation (e.g., by running serial dictatorship). The work of Manshadi et al. [34] is conceptually similar to ours but addresses a different problem, where agents arrive dynamically and request quantities of a divisible good.

1.1 Why should we care about order symmetry?

Fairness is a rather nebulous concept, and is often thought of in terms of outcomes after preferences are known (*outcome fairness* or *equity*). We cannot achieve complete outcome fairness with a deterministic discrete mechanism in the worst case (as shown by a unanimous profile). Instead, order symmetry guarantees average-case fairness. Order symmetry can also be thought of as outcome fairness where the uncertainty of the mechanism designer about agent preferences is encoded by the probability measure *P*, and fairness in expectation over *P* is the desired result.

We now give two more related justifications for order symmetry. A competing idea of fairness is *procedural fairness* or *equality of opportunity*, which requires *a priori* symmetry between agents, with outcomes determined solely by their preferences. In social choice theory, this is typically formalized as anonymity, but anonymity is unattainable for discrete assignment mechanisms. A natural relaxation is order symmetry, which ensures that, absent preference information, no agent has a systematic advantage. Order symmetry also relates to *accountability* and *transparency* on the part of the mechanism designer. If all agents receive the same expected outcome quality and the mechanism is chosen before preferences are known, the designer is insulated from manipulation claims, as no agent can be placed in a systematically stronger or weaker position.

As anecdotal evidence that order symmetry aligns with human values, we conducted a short poll of researchers attending the 2023 COMSOC workshop on computational social choice. We simply defined for them the standard and reversing forms of the Boston mechanism (see Section 5) and asked which of the two they perceived to be more fair. Out of 42 respondents, 24 (57%) said that the reversing variant was more fair, 17 (41%) said that both were equally fair, and only 1 (2%) said that the standard variant was more fair. Furthermore several participants provided justifications that alluded to equalizing outcomes for agents having different positions. As we will see in Section 5, the two versions achieve the same welfare but the reversing variant is less biased than the standard variant (this was not told to the participants in the poll).

1.2 Our contribution, and outline of the paper

Our main conceptual contribution is to define the concept of *order symmetry* (Section 3), a natural and novel fairness guarantee for assignment mechanisms, defined with respect to a probability measure on profiles. We additionally define *order bias*, a quantitative measure of the failure of order symmetry, which allows us to compare mechanisms even in settings where perfect order symmetry cannot or ought not be satisfied. A common thread throughout the paper is the use of order symmetry/bias as a lens for mechanism design—we see on several occasions that variants of common mechanisms can be defined to improve fairness without sacrificing other guarantees. Our claims are substantiated through a combination of axiomatic, probabilistic, and numerical analysis.

In Section 4, we compare and contrast SD and TTC from the perspective of order symmetry. Although both mechanisms are strategyproof and efficient, we show that TTC is order-symmetric with respect to any fully symmetric probability measure while SD has the maximum order bias among a natural class of mechanisms with respect to the uniform measure. Using the perspective of order symmetry, we consider two variants of SD, including one which preserves strategyproofness and efficiency, and outperforms TTC in terms of order bias in numerical simulations for a wide range of Mallows distributions.

In Section 5, we examine the common-tiebreak Boston mechanism, which is efficient but not strategyproof. Though not generally order-symmetric, except for very specific probability measures, it exhibits low order bias in simulations. To further reduce bias, we propose two independent improvements. The first introduces cyclic, object-specific priorities over the agents, making the mechanism order-symmetric under fully symmetric probability measures and yielding strong benefits in simulation. (Technically this mechanism is a member of the class of Boston mechanisms, with the standard variant having all objects use the same priority ordering over the agents.) The second modifies object priorities via reversal during the algorithm, creating novel mechanisms outside the Boston class and achieving a moderate but meaningful reduction in order bias.

2 PRELIMINARIES

Let *n* be a positive integer and let $O = \{o_1, ..., o_n\}$ be an ordered set of **objects** having cardinality *n* and $A = \{a_1, ..., a_n\}$ 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 $\pi = (\succ_1, ..., \succ_n) = (\pi(1), ..., \pi(n))$. We write the preference order of agent *k* as $(\pi(k)_1, ..., \pi(k)_n)$. We let $\Pi(A, O)$ denote the set of all preference profiles.

An **assignment** is a bijective mapping taking *A* to *O*. Note that each assignment can be represented by a permutation matrix with rows indexed by agents and columns by objects. For an assignment α we write $\alpha(i)$ to denote the object assigned to agent *i*. An **assignment mechanism** \mathcal{A} is a function associating an assignment $\mathcal{A}(\pi)$ with each profile π .

As is common, we assume agents care only about their assigned object, allowing us to infer their preferences over assignments. This may include indifferences, as an agent might receive the same object in multiple assignments. In order to simplify notation, we can write $\alpha \geq_i \beta$ to mean that agent *i* weakly prefers assignment α to β , and $\alpha >_i \beta$ to mean a strict preference. In terms of our previous notation, $\alpha >_i \beta$ if and only if $\alpha(i) >_i \beta(i)$.

It will be helpful to introduce two common assignment mechanisms. The **Serial Dictatorship** (SD) mechanism works as follows: fix an exogenous order ρ on the agents, and let them choose in turn according to ρ their highest ranked object from those remaining. Let us also define the **Top Trading Cycles** (TTC) algorithm, originally presented by Shapley and Scarf [42] and attributed to David Gale. Fix an initial assignment (or endowment) of objects to agents. Each agent *i* points to the agent currently assigned their most preferred object. It is clear that there must be at least one cycle (including possible self-cycles if an agent already holds their top choice). For each cycle, execute the indicated trade by reassigning each object to

the agent pointing to it, then remove all involved agents and objects. Repeat this process—having agents point to their most preferred remaining object and clearing cycles—until no cycles remain. After all cycles have been cleared, have the agents once again point to their most-preferred remaining object and continue clearing cycles; terminate when there are no cycles left.

Example 2.1. Consider the profile where agents a_1, a_2, a_3 have respective preferences over objects o_1, o_2, o_3 as follows.

```
a_1: o_1 > o_2 > o_3, \quad a_2: o_1 > o_3 > o_2, \quad a_3: o_2 > o_1 > o_3.
```

SD with picking order a_1, a_2, a_3 assigns o_1 to a_1, o_3 to a_2 , and o_2 to a_3 . It is easy to check that the same assignment is output by running TTC with initial endowment that assigns object o_i to agent a_i for each *i*. On the other hand, if the initial endowment had assigned $a_1 \leftarrow o_2, a_2 \leftarrow o_1, a_3 \leftarrow o_3$ then the resulting assignment gives o_2 to a_1, o_1 to a_2 , and o_3 to a_3 .

We conclude this section with some important properties of assignment mechanisms. Efficiency requires that an assignment mechanism never outputs an assignment for which some agent can be made better off without hurting any other agent.

Definition 2.2 (efficiency). Let α be an assignment. We say that α is **efficient** if there does not exist an assignment α' such that $\alpha'(k) \geq_i \alpha(k)$ for all agents $k \in A$ and $\alpha'(i) >_i \alpha(i)$ for some agent $i \in A$. An assignment mechanism \mathcal{A} is called efficient if $\mathcal{A}(\pi)$ is efficient for all $\pi \in \Pi(A, O)$.

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

Definition 2.3 (strategyproofness). Assignment mechanism \mathcal{A} is **strategyproof** if, for all profiles $\pi \in \Pi(A, O)$ and all reports \succ'_i , it holds that $\mathcal{A}(\pi) \geq_i \mathcal{A}(\succ'_i, \pi_{-i})$.

SD and TTC are both known to be efficient and strategyproof.

Finally, we define anonymity, a common fairness criterion which says that assignments should depend only on the preferences of the agents, and not on their identities. We define it in terms of group actions, which will be a helpful framework for the remainder of the paper. Let *G* denote the group Sym(*A*) of all permutations of the set of agents. Since we have ordered the set $A = \{a_1, \ldots, a_n\}$ we may identify *G* 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 $(g \cdot \alpha)(a) = \alpha(g^{-1}(a))$. As argued earlier, no discrete assignment mechanism is anonymous.

Definition 2.4 (anonymity). Assignment mechanism \mathcal{A} is **anonymous** if and only if $\mathcal{A}(g \cdot \pi) = g \cdot \mathcal{A}(\pi)$ for each $g \in G, \pi \in \Pi(A, O)$.

3 ORDER SYMMETRY

We now formalize the definition of order symmetry.

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 $\{1,\ldots,n\} \times \{1,\ldots,n\}$ whose value at (r, j) is 1 if \mathcal{A} assigns agent r the object it ranks as jth best and 0 otherwise.

As an illustration, consider the instance from Example 2.1. 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 its second-top object.

Definition 3.2. Let *P* be a probability measure on the set of all profiles for a given *n*. The assignment mechanism \mathcal{A} is **order**-symmetric with respect to *P* if for all *j*, the following quantity is independent of *r*:

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

Thus a mechanism is order-symmetric if and only if all rows of the rank distribution matrix are equal in expectation. The following example illustrates the idea. *Impartial Culture* (IC) is the uniform probability measure on the set of all profiles.

Example 3.3. Consider the case of agents a_1, a_2, a_3 and objects o_1, o_2, o_3 , 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 $o_1 > o_2 > o_3$. 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 its 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 $\begin{bmatrix} 1 & 0 & 0 \\ 2/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}$. For TTC with initial endowment that assigns o_i to a_i for $i \in \{1, 2, 3\}$, the expected rank distribution matrix under IC follows from computing the outcome on each profile and is $\begin{bmatrix} 2/3 & 2/9 & 1/9 \\ 2/3 & 2/9 & 1/9 \\ 2/3 & 2/9 & 1/9 \\ 2/3 & 2/9 & 1/9 \end{bmatrix}$.

From Examples 2.1 and 3.3, we see that TTC is order-symmetric with respect to IC for n = 3, but not with respect to the point mass concentrated at the profile from Example 2.1. Also, SD with picking order a_1 , a_2 , a_3 is not order-symmetric with respect to either measure. Given these observations, it is natural to ask which probability measures allow for the existence of order-symmetric mechanisms. As the following example shows, there exist measures for which no mechanism is order-symmetric, and measures for which order symmetry is an easy condition to satisfy.

Example 3.4. Consider a unanimous profile, and let *P* put all its weight on this profile. Then every discrete assignment mechanism allocates exactly one agent its most-preferred object, so that the mechanism cannot be order-symmetric with respect to *P*. On the other hand, consider a measure with all its weight on *contention-free* profiles, i.e., those where all agents have a different first preference [43]. Any efficient discrete assignment mechanism will assign each agent their top choice, thereby satisfying order symmetry.

3.2 Order bias

Example 3.4 shows that not all probability measures permit existence of an order-symmetric mechanism. Even for more permissive measures, not all interesting assignment mechanisms are ordersymmetric. It is therefore useful to measure the deviation of a mechanism from order symmetry. To define a quantitative measure of bias with access to only ordinal information, we force a common utility function on agents via a scoring rule.

Definition 3.5. A **positional scoring rule** is given by a sequence s of real numbers $s_1 \ge s_2 \ge \cdots \ge s_n$ with $s_1 > s_n$.

Commonly used scoring rules include *plurality* defined by (1, 0, 0, ..., 0), *antiplurality* defined by (1, 1, 1, ..., 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 3.6. The *expected utility* of agent *r* with respect to a scoring rule *s*, mechanism \mathcal{A} and measure *P* on profiles is

$$U(r) := U_{\mathcal{A},s,P}(r) = \sum_{j} s_{j} E_{P}[D_{\pi,\mathcal{A}}(r,j)].$$

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 mostand least-preferred objects:

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

With the given normalization, $0 \le \beta_n(\mathcal{A}, s, P) \le 1$. If \mathcal{A} is ordersymmetric then $\beta_n(\mathcal{A}, s, P) = 0$, 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. While the latter situation will not happen often in practice, it can happen, for example when P has all its weight on profiles in which all agents have the same preference order, and SD is used.

Example 3.7. Let $1 \le k \le n$ and let P_k be any probability measure whose support consists entirely of profiles for which agents $1, \ldots, k$ have the same preference, but the last n - k agents have different first choices taken from the n - k bottom choices of the first k agents. On each profile in the support of P_k , SD with choosing order $1, 2, \ldots, n$ assigns agent i its ith choice for $1 \le i \le k$ and the other agents their first choice. The same is true of TTC 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 $(s_1 - s_k)/(s_1 - s_n)$.

In the following sections, we use order symmetry and order bias to discriminate between mechanisms on fairness grounds.

4 TOP TRADING CYCLES AND SERIAL DICTATORSHIP

In this section we focus on TTC and SD, two of the most common mechanisms in practical use. Both are efficient and strategyproof. However, as we have alluded to in our illustrative example with Alice and Bob (Section 1), SD is defined in a way that gives an advantage to agents early in the picking order, whereas TTC does not incorporate an inherent advantage for any agents, in the absence of knowledge about the preferences. We make this intuition precise in this section using our notion of order symmetry.

4.1 Fully symmetric measures and TTC

Considering Example 3.4, we would like to find a class of measures general enough to capture real-world settings, but not so general as to preclude the existence of order-symmetric mechanisms. We introduce such a class in this subsection. Recall from our definition of anonymity that *G* denotes the group Sym(\mathcal{A}) of all permutations of the set of agents. Similarly, let *H* denote the group Sym(\mathcal{O}) of all permutations of the set of objects. The group *G* acts on $\Pi(\mathcal{A}, \mathcal{O})$ via $(g \cdot \pi)(a) = \pi(g^{-1}(a))$, and on assignments by $(g \cdot \alpha)(a) = \alpha(g^{-1}(a))$, while *H* acts on $\Pi(\mathcal{A}, \mathcal{O})$ by $(h \star \pi)(a) = (h(\pi(a)_1), \cdots, h(\pi(a)_n)$ and on assignments by $(h \star \alpha)(a) = h(\alpha(a))$. 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 $g \cdot (h \star \pi) = h \star (g \cdot \pi)$ for all $q \in G, h \in H$.

Definition 4.1. An **anonymous probability measure** on Π is one for which $P(g \cdot \pi) = P(\pi)$ for all $g \in G$ and all $\pi \in \Pi$. A **neutral probability measure** on $\Pi(A, O)$ is one for which $P(h \star \pi) = P(\pi)$ for all $h \in H$ and all $\pi \in \Pi$. A probability measure is **fully symmetric** if it is anonymous and neutral. In other words, $P(g \cdot (h \star \pi)) = P(\pi)$ for all $g \in G, h \in H$.

Note that *P* is fully symmetric if and only if *P* takes a constant value on each orbit of $G \times H$.

Example 4.2. Examples of fully symmetric measures include Impartial Culture (each profile has the same weight), Impartial Anonymous Culture (any two profiles in the same orbit under G have the same weight and every orbit has the same total weight), Impartial Anonymous Neutral Culture [19], and indeed the uniform measure on any set of profiles defined without reference to an order of agents or objects. For example, the set of all single-peaked profiles is a common model for preferences where there is a single dimension on which to evaluate objects.

A measure concentrated on a single unanimous profile is anonymous but not neutral. The uniform measure on profiles where all even-indexed agents have the same preference and all odd-indexed agents have the same preference is neutral but not anonymous.

We are now able to show our first main result, that TTC is order-symmetric with respect to any fully symmetric measure. The result generalizes Example 3.3, which showed that TTC is order-symmetric with respect to IC for n = 3. Intuitively, the result is true because, under the assumption of symmetry between agents and objects, no agent is systematically advantaged by their initial endowment. The proof of Theorem 4.3, along with all other proofs, can be found in the full version of the paper.

THEOREM 4.3 (TTC IS ORDER-SYMMETRIC). TTC with any fixed endowment is order-symmetric with respect to every fully symmetric probability measure.

One may wonder whether Theorem 4.3 characterizes fully symmetric profiles, but it does not. Consider a measure P that puts all its weight on a single contention-free profile. Then TTC is order-symmetric according to P even though P is not fully symmetric.

4.2 SD is far from order-symmetric

We have seen already in Example 3.3 that SD is not order-symmetric with respect to IC. It is easy to guess that under SD, every row of the expected rank distribution matrix stochastically dominates each one below it, so that the last agent is the worst off. We can give much more precise results, and compute the rank distribution matrix under IC exactly for each n.

THEOREM 4.4. Let \mathcal{A} be SD with picking order equal to $\{a_1, \ldots, a_n\}$, and let P be IC. Then

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

In particular if r = n, then agent r is equally likely to get each possible object. Combined with the fact that the first agent always attains its first choice under SD, we have the following corollary.

COROLLARY 4.5. We have $\beta_n(SD, s, IC) = \frac{s_1 - \overline{s}}{s_1 - s_n}$, where \overline{s} is the mean of s.

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

Proposition 4.6 says that SD is in fact the most order-biased of a wide class of mechanisms. For each measure P, we define C(P) to be the class of mechanisms for which the expected rank distribution of each agent stochastically dominates the uniform distribution. This is a weak welfare requirement that rules out, for example, the mechanism that reverses the input profile before running SD.

PROPOSITION 4.6. For every scoring rule, SD attains the maximum order bias in C(IC).

Theorem 4.3 and Proposition 4.6 are striking in that they maximally separate TTC and SD in terms of order bias. These two mechanisms, the two most prominent both in theory and practical use, are both efficient and strategyproof, and their randomized versions are known to be equivalent in a strong sense [1, see Section 6 for additional discussion]. However, as exemplified by the Alice and Bob example in the introduction, the two mechanisms feel quite different in terms of fairness properties, and order bias turns out to be the right definition to capture that difference.

4.3 More symmetric variants of SD

Despite its high order bias, SD does have many qualities that explain its widespread use in practice: it is simple to understand and play, it requires each agent to report a choice of only a single object (as opposed to a complete ranking), and it is strategyproof and efficient. One may wonder whether these qualities can be preserved while also mitigating its poor order bias.

One idea is to fix a bijection between agents and objects. For simplicity, suppose agent a_i is mapped to object o_i for all *i*. Fix an ordering of the agents, without loss of generality (a_1, a_2, \ldots, a_n) . Define SD_c as the mechanism in which the first *c* agents choose an object one at a time according to the fixed ordering as in SD, but the remaining n - c agents choose in order determined by the relative position of their corresponding object in agent a_1 's preference. Agents who map to more-preferred objects choose earlier than agents who map to less-preferred objects. Note that SD_n and SD_{n-1} are simply SD, while SD₀ has the picking order completely determined by the reported preference of a_1 .

Example 4.7. Let n = 3 and suppose that agent a_1 reports preference $o_3 > o_1 > o_2$. Under SD₀, the picking order is (a_3, a_1, a_2) . Under SD₁, the picking order is (a_1, a_3, a_2) , since a_1 ranks o_3 above o_2 . Finally, SD₂ and SD₃ result in picking order (a_1, a_2, a_3) .

 SD_c mechanisms are members of the class of *sequential choice rules* [38], a class which also includes sequential dictatorships [e.g., 37]. Importantly, note that the picking orders for SD_c are allowed to vary by profile, unlike in SD. We are not aware of a formal treatment of SD_c mechanisms in the literature, but note that they have been used to illustrate bossiness [e.g., 45, Footnote 19].²

In contrast to SD, SD_c requires a_1 to report a complete preference order, but, as in SD, all other agents can participate in the mechanism by simply choosing an object from the available set when their turn comes. It is easy to see that SD₀ is order-symmetric with respect to IC because the randomness in the preference of agent a_1 is being inherited by the mechanism itself in such a way that all picking orders are equally likely. However, strategyproofness is violated by allowing a_1 so much control over the picking order.

THEOREM 4.8. SD_0 is order-symmetric with respect to any fully symmetric probability measure but violates strategyproofness.

Importantly, order symmetry of SD_0 relies on the assumption that agent a_1 will truthfully report its preference. We often might not expect the mechanism to be order-symmetric in practice even with respect to preferences that are governed by the IC measure, since agent a_1 can manipulate the outcome in their favor. Instead, we can consider SD_1 , which guarantees strategyproofness while providing order bias intermediate between SD and SD₀.

THEOREM 4.9. SD₁ is strategyproof.

The next theorem gives the order bias of SD₁ with respect to IC for any scoring rule. The proof is a probabilistic analysis that reveals an interesting fact about the SD₁ mechanism: All agents a_j , $j \ge 2$, get their first-choice object with probability exactly $\frac{1}{2}$.

THEOREM 4.10. SD₁ has order bias $\frac{s_1-s^*}{s_1-s_n}$ with respect to IC, where

$$s^* = \frac{1}{2}s_1 + \sum_{j=2}^n s_j \frac{1}{n-1} \sum_{r=j}^n \frac{\binom{n-j}{r-j}}{\binom{n}{r-1}}$$

The order bias of SD₁ under IC is 1/2 for plurality and $1 - \frac{1}{n(n+1)}$ for antiplurality. The order bias with respect to Borda is $O(\frac{\log n}{n})$ as $n \to \infty$, and hence has limiting value zero.

For Borda, the order bias for SD₁ is much smaller than the value 1/2 for SD even for small values of *n*. For example, when n = 4 it is approximately 0.26852, while when n = 32 it is approximately 0.07277. Under IC, all agents a_2, \ldots, a_n receive the same expected utility, as the mechanism treats them symmetrically. Only agent a_1 receives a higher expected utility by virtue of its privileged position.

4.4 Numerical simulation results

In the last subsection, we derived analytic results for order bias under IC, which is a helpful and tractable baseline, but often does not reflect realistic preference structures. In real-world assignment problems, some objects are typically more highly sought after than others, but this is not captured by IC. To explore a more realistic preference model, we resort to simulation and consider the Mallows model [33]. The Mallows model is parametrized by a *reference* order σ and a dispersion parameter $\phi \in (0, 1]$. Let > be a strict linear order. Then the Mallows model specifies that the probability an agent has preference order > is $P(>) = P(> |\sigma, \phi) = \frac{1}{Z}\phi^{d(>,\sigma)}$, where Z is a normalization parameter and $d(>, \sigma)$ is the number of pairwise disagreements between > and σ . 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 , and that when $\phi = 1$ the distribution is maximally dispersed, yielding Impartial Culture. Unless $\phi = 1$ the Mallows measure is not neutral, but it is always anonymous.

Figure 1 shows order bias and utilitarian welfare³ for Borda utilities⁴ as we vary the Mallows parameter across its whole range for fixed *n*. We chose n = 32 as a representative value that is not "too small" and yet is computationally tractable. Note that, as expected, all algorithms have order bias equal to 1 when the Mallows parameter is 0, indicating full concentration on a single unanimous preference profile. For all higher values of ϕ , SD(=SD_n) has the highest order bias. In terms of order bias, SD₀ performs best across all Mallows parameters, but recall that this comes at the price of violating strategyproofness, while TTC outperforms SD for all parameter values. A tantalizing question that we leave open is whether this relationship holds in general for all anonymous measures. Note that SD1 achieves lower order bias than TTC for all except very high values of ϕ , suggesting that it may be a robust choice when both order bias and strategyproofness concerns are paramount. Finally, we observe that the utilitarian welfare of all four algorithms is identical, suggesting that any benefits in terms of order bias do not come at a cost in terms of welfare.



Figure 1: (left) Borda order bias, n = 32, (right) Borda utilitarian welfare, n = 32. Mean of 5000 simulations.

5 BOSTON MECHANISMS

For fully symmetric probability measures, we have seen that TTC is a strong candidate for practical use – it is efficient, strategyproof, and order-symmetric. However, for Mallows-distributed preferences with even a medium concentration parameter, TTC has only

²By definition, a bossy mechanism is one in which some agent can change the allocation for other agents without changing their own allocation.

³Utilitarian welfare for mechanism \mathcal{A} with respect to measure P and scoring rule s is the expectation under P of the arithmetic mean score received by an agent. Simply put, it is the expected average utility.

⁴We use Borda utilities throughout the main text as they are a common and natural choice. In the full version of the paper, we replicate all results for 3-approval utilities. Note that *k*-approval utilities are commonly used to evaluate assignments in practice (in particular, in school choice settings), as they place emphasis on an agent receiving one of her top *k* choices, a simple metric to report and understand.

slightly lower order bias than SD, and we might expect this behavior to hold whenever there is a sense in which the initial endowment favors some agents more than others (e.g., there is a predictable ranking of the objects). In this section we explore the Boston mechanisms, an alternative class of mechanisms that show strong performance on Mallows-distributed preferences, and we apply the order symmetry lens to design novel modifications of Boston.

5.1 Definition and background

The **Boston mechanism** (henceforth **Boston**), named for its use to match students to schools in Boston until 2005, is commonly used in school choice settings. For example, variants of Boston are used in Seattle WA, Charlotte NC, Barcelona [14], and across Germany [8], among other places. Here we consider the restriction to the housing allocation model, so that each "school" is a single object.

The mechanism 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 an object-dependent priority ordering ρ_j is used as a tiebreak to allocate it to one of them.⁵ Since each priority profile induces a Boston mechanism, the definition in fact yields a class of mechanisms, which were axiomatically characterized by Kojima and Ünver [28]. A natural and common special case is the version where ρ_j is independent of the object *j*, which corresponds to a single tiebreak order being applied to all conflicts. We will refer to this version as the common-tiebreak variant of Boston.

5.2 Boston and Order Symmetry

In this section we analyze the Boston mechanisms from the perspective of order symmetry and order bias. Our goal is twofold. First, we want to find a mechanism that is more robust to non-neutral probability measures than TTC, in terms of order bias. Second, we aim to provide guidance as to which member of the class of Boston mechanisms to use in practice.

Rather than using a common tiebreak for all objects, which induces asymmetry in favor of agents higher in the tiebreak order, order symmetry considerations point to using a more balanced priority profile. In particular, we will consider the cyclic priority profile defined by the priority order $\rho_j = a_j > a_{j+1} > ... > a_n > a_1 > ... > a_{j-1}$ for object o_j . Of course, agent a_j has an advantage over other agents with respect to object o_j , so a_j is advantaged to the extent that o_j is systematically preferred to other objects. However, when the preference distribution is neutral and anonymous, cyclic-tiebreak Boston is order-symmetric. As was the case for TTC, the intuition is that no agent is systematically favored by the priority profile provided that we cannot distinguish between agents and objects a priori.

THEOREM 5.1. Cyclic-tiebreak Boston is order-symmetric with respect to any fully symmetric probability measure.

Numerical simulation results. We compare cyclic-tiebreak Boston (Bcyc) and common-tiebreak Boston (Bcom) on Mallows data in Figure 2. For both Boston variants, the order bias drops sharply as the Mallows parameter increases from 0. The drop is particularly

sharp for the cyclic-tiebreak variant, which reaches an order bias very close to 0.5 for $\phi \approx 0.1$. Recall that, in contrast, SD only achieves order bias of 0.5 on the easy case of IC preferences. The order bias continues to decrease with ϕ , reaching 0 and 0.13 at $\phi = 1$ for cyclic- and common-tiebreak Boston respectively. Once again, we see that utilitarian welfare is identical between the two variants, indicating that the order bias improvement from the cyclic-tiebreak variant is obtained without any cost in welfare.

In the same figure we show the results for TTC for comparison. Both Boston variants, but especially cyclic-tiebreak Boston, exhibit a striking dominance both in terms of order bias and welfare. It is particularly interesting to note the shape of the lines in the order bias plot. Cyclic-tiebreak Boston and TTC are both order-symmetric for $\phi = 1$, but Boston's low order bias is fairly robust across a wide spectrum of parameters, whereas TTC is much more fragile. For intuition, when some objects are systematically preferred to others then TTC is biased in favor of agents who receive favorable objects in the preliminary assignment. The preliminary assignment is "sticky" in the sense that no agent is worse off in the TTC assignment than in the preliminary assignment. Accordingly, the TTC allocation is quite constrained. On the other hand, the Boston mechanisms maximize the number of agents who receive their first preference, then their second, and so on, without protecting any specific agent from receiving a lowly-preferred object.



Figure 2: (left) Borda order bias, n = 32, (right) Borda utilitarian welfare, n = 32. Mean of 5000 simulations.

5.3 Reversing Boston for lower order bias

We have seen that a cyclic priority profile achieves significantly lower order bias than the common-tiebreak Boston variant, as well as TTC and SD. Can we do better still? In this section we provide one technique for doing so. For intuition, consider the commontiebreak Boston mechanism. Being last in the tiebreak order is a substantial disadvantage relative to being first, since conflicts are never resolved in favor of the last agent. Intuitively, using the same tiebreak order in every round places a large importance on whatever process is used to generate the order, which is antithetical to the idea of order symmetry. The modification that we consider here is to reverse the tiebreak order after every round, so that odd-numbered rounds use the original tiebreak order while even-numbered rounds use the reversed order.⁶ To generalize this idea to other Boston mechanisms, we simply reverse the priority order of all objects in

⁵Mennle and Seuken [35] discuss a variant of Boston that they call *Adaptive Boston*, but in the interest of concise presentation we do not consider it here.

⁶One could imagine various other reversing patterns as well as modifications to the tiebreak order beyond reversals, but we do not consider such variants explicitly.

even-numbered rounds. We call the resulting algorithms *reversing Boston mechanisms*. In the full version of the paper we provide an example to show that reversing Boston mechanisms cannot be written as Boston mechanisms in general; to our knowledge, these mechanisms have not been proposed before.

Numerical simulation results. Figure 3 compares the performance of the ordinary and reversing forms of common-tiebreak Boston (Bcom and Rcom respectively), along with the ordinary and reversing forms of cyclic-tiebreak Boston (Bcyc and Rcyc respectively). Both reversing forms have lower order bias than the ordinary forms, with the reversing form of common-tiebreak Boston significantly outperforming the original. The reversing form of cyclic-tiebreak Boston exhibits a modest but non-trivial decrease in order bias for Mallows parameter between 0.1 and 0.5.



Figure 3: (left) Borda order bias, n = 32, (right) Borda utilitarian welfare, n = 32. Mean of 5000 simulations.

6 **DISCUSSION**

We introduced and analyzed order symmetry, a natural fairness concept ensuring all agents perform equally well in expectation. With a sufficiently symmetric probability measure, order symmetry serves as an average-case analog of anonymity-one that is achievable by deterministic assignment algorithms without sacrificing other desirable properties. When exact order symmetry is unattainable, some order bias is inevitable. Our analysis suggests that when low order bias is desired alongside strategyproofness, SD should be replaced by SD₁ or TTC. Where strategyproofness is not essential, the common-tiebreak Boston mechanism performs well in terms of order bias, significantly outperforming the strategyproof mechanisms on Mallows data for all but very large values of dispersion parameter. It can be improved even further by employing the cyclic tiebreak scheme and/or reversing the tiebreak order after the first round. We conclude with some additional observations and directions for future work.

Randomized mechanisms

This paper focuses on order symmetry in deterministic mechanisms, but many practical house allocation and school choice mechanisms are randomized. Typically, a deterministic mechanism is randomly selected and implemented—for example, by randomly choosing a selection order (for SD) or a priority profile (for Boston) while respecting fixed priorities like seniority or location. For TTC, randomization involves selecting an initial endowment. Our results have direct implications for randomized mechanism design. In a famous result, Abdulkadiroğlu and Sönmez [1] showed that uniform randomized versions of SD and TTC are exactly equivalent in terms of the lotteries over assignments that they output. Note that these randomized versions are anonymous (at the stage before the randomization is realized). However, participants might also care about fairness after the randomization has produced a deterministic mechanism, and randomizing over order-symmetric (or, more generally, low order bias) deterministic mechanisms provides a fairness guarantee. While a formal treatment of order symmetry of randomized mechanisms is beyond the scope of this paper, it is easy to see that in this sense, the randomized versions of SD and TTC are very different.

Multi-unit assignment

We have not examined cases where the number of objects exceeds the number of agents and each agent receives a bundle of objects, but order bias remains relevant and could be generalized by extending preferences from individual objects to sets of objects [7]. Many mechanisms, such as sports league drafts, use a "picking sequence" that significantly impacts outcomes. Various methods have been proposed to balance outcomes between agent positions [9, 11, 12, 26, 27], and order bias could serve as a useful tool for comparing these 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 [30], envy-freeness up to any object [15], and proportionality up to one object [16] have been widely studied, and prior work has considered combining fairness guarantees both before and after randomization in more general multi-object assignment settings [3–5, 13]. None of this work is relevant to the house allocation setting since the fairness guarantees are too permissive when each agent receives only a single object.

General probability measures.

A compelling direction for future work is to dive more deeply into the possibilities and limitations for general probability measures. For practical application, we imagine order symmetry being incorporated into mechanism design in two steps. First, the designer estimates (say, from historical data) the probability measure from which the preferences will be drawn. Second, the designer searches 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.

ACKNOWLEDGMENTS

We thank Michael Albert and Jens Witkowksi for helpful comments on preliminary versions of the paper, and numerous anonymous journal and conference referees for their feedback.

REFERENCES

- Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.
- [2] David J. Abraham, Katarína Cechlárová, David F. Manlove, and Kurt Mehlhorn. 2004. Pareto optimality in house allocation problems. In *International Symposium on Algorithms and Computation*. Springer, 3–15.
- [3] Mohammad Akbarpour and Afshin Nikzad. 2020. Approximate random allocation mechanisms. *Review of Economic Studies* 87, 6 (2020), 2473-2510.
- [4] Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2024. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research* 72, 4 (2024), 1674–1688.
- [5] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2022. On best-of-both-worlds fair-share allocations. In *International Conference on Web and Internet Economics*. Springer, 237–255.
- [6] Sophie Bade. 2020. Random serial dictatorship: the one and only. Mathematics of Operations Research 45, 1 (2020), 353–368.
- [7] Salvador Barberà, Walter Bossert, and Prasanta K. Pattanaik. 2004. Ranking sets of objects. In *Handbook of utility theory*. Springer, 893–977.
- [8] Christian Basteck, Katharina Huesmann, and Heinrich Nax. 2015. Matching Practices for Secondary Schools—Germany. *MiP Country Profile* 21 (2015).
- [9] Aurélie Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, and Parham Shams. 2019. Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods. In Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems. 900–908.
- [10] Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. Journal of Economic Theory 100, 2 (2001), 295–328.
- [11] Sylvain Bouveret and Jérôme Lang. 2011. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*. 73–78.
- [12] Steven J. Brams and Alan D. Taylor. 2000. The win-win solution: Guaranteeing fair shares to everybody. WW Norton & Company.
- [13] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. Designing random allocation mechanisms: Theory and applications. *American Economic Review* 103, 2 (2013), 585–623.
- [14] Caterina Calsamiglia and Maia Güell. 2018. Priorities in school choice: The case of the Boston mechanism in Barcelona. *Journal of Public Economics* 163 (2018), 20–36.
- [15] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. ACM Transactions on Economics and Computation 7, 3 (2019), 1–32.
- [16] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair public decision making. In Proceedings of the 2017 ACM Conference on Economics and Computation. 629–646.
- [17] Hervé Crès and Hervé Moulin. 2001. Scheduling with opting out: Improving upon random priority. *Operations Research* 49, 4 (2001), 565–577.
- [18] Xiaotie Deng, Yansong Gao, and Jie Zhang. 2017. Smoothed and Average-Case Approximation Ratios of Mechanisms: Beyond the Worst-Case Analysis. In Proceedings of the 42nd International Symposium on Mathematical Foundations of Computer Science, Vol. 83. 16.
- [19] Ömer Eğecioğlu and Ayça E. Giritligil. 2013. The impartial, anonymous, and neutral culture model: a probability model for sampling public preference structures. *Journal of Mathematical Sociology* 37, 4 (2013), 203–222.
- [20] Haluk I. Ergin. 2000. Consistency in house allocation problems. Journal of Mathematical Economics 34, 1 (2000), 77–97.
- [21] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. Social Welfare in One-Sided Matchings: Random Priority and Beyond. In *Proceedings*

of the 7th International Symposium on Algorithmic Game Theory. 1–12.

- [22] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. 2019. Envyfreeness in house allocation problems. *Mathematical Social Sciences* 101 (2019), 104–106.
- [23] Yansong Gao and Jie Zhang. 2019. Average-case analysis of the assignment problem with independent preferences. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 287–293.
- [24] Patrick Harless and Vikram Manjunath. 2018. Learning matters: Reappraising object allocation rules when agents strategically investigate. *International Economic Review* 59, 2 (2018), 557–592.
- [25] Aanund Hylland and Richard Zeckhauser. 1979. The efficient allocation of individuals to positions. *Journal of Political Economy* (1979), 293–314.
- [26] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. 2013. A social welfare optimal sequential allocation procedure. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*. 227–233.
- [27] David A. Kohler and R. Chandrasekaran. 1971. A class of sequential games. Operations Research 19, 2 (1971), 270–277.
- [28] Fuhito Kojima and M Utku Ünver. 2014. The "Boston" school-choice mechanism: an axiomatic approach. *Economic Theory* 55 (2014), 515–544.
 [29] Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. 2019. Size
- [29] Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. 2019. Size versus truthfulness in the house allocation problem. *Algorithmica* 81, 9 (2019), 3422–3463.
- [30] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the* 5th ACM Conference on Electronic Commerce. 125–131.
- [31] Xinghua Long. 2016. Essays on house allocation problems. Ph. D. Dissertation. Department of Economics, Texas A&M University. https://hdl.handle.net/1969. 1/157810
- [32] Xinghua Long and Rodrigo A. Velez. 2021. Balanced House Allocation. arXiv preprint arXiv:2109.01992 (2021).
- [33] Colin L. Mallows. 1957. Non-null ranking models. I. Biometrika 44, 1/2 (1957), 114–130.
- [34] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. 2023. Fair Dynamic Rationing. *Management Science* 69, 11 (2023), 6818–6836.
- [35] Timo Mennle and Sven Seuken. 2017. Trade-offs in school choice: comparing deferred acceptance, the classic and the adaptive Boston mechanism. arXiv preprint arXiv:1406.3327 (2017).
- [36] Ali I Ozkes and M Remzi Sanver. 2021. Anonymous, neutral, and resolute social choice revisited. *Social Choice and Welfare* 57, 1 (2021), 97–113.
- [37] Szilvia Pápai. 2000. Strategyproof assignment by hierarchical exchange. Econometrica 68, 6 (2000), 1403–1433.
- [38] Marek Pycia and Peter Troyan. 2023. A theory of simplicity in games and mechanism design. *Econometrica* 91, 4 (2023), 1495–1526.
- [39] Marek Pycia and M. Utku Unver. 2014. Incentive Compatible Allocation and Exchange of Discrete Resources. (2014). Working paper.
- [40] Daniela Saban and Jay Sethuraman. 2013. House allocation with indifferences: a generalization and a unified view. In *Proceedings of the 14th ACM Conference* on *Electronic Commerce*. 803–820.
- [41] Daniela Saban and Jay Sethuraman. 2015. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research* 40, 4 (2015), 1005–1014.
- [42] Lloyd Shapley and Herbert Scarf. 1974. On cores and indivisibility. Journal of Mathematical Economics 1, 1 (1974), 23–37.
- [43] Priyanka Shende and Manish Purohit. 2023. Strategy-proof and envy-free mechanisms for house allocation. *Journal of Economic Theory* 213 (2023), 105712.
- [44] Tayfun Sönmez and M. Utku Ünver. 2010. House allocation with existing tenants: A characterization. Games and Economic Behavior 69, 2 (2010), 425–445.
- [45] William Thomson. 2016. Non-bossiness. Social Choice and Welfare 47, 3 (2016), 665–696.