Majoritarian Assignment Rules

Felix Brandt
Technical University of Munich
Munich, Germany
brandt@tum.de

Haoyuan Chen
Technical University of Munich
Munich, Germany
ge85hud@tum.de

Chris Dong
HPI, University of Potsdam
Potsdam, Germany
chrisshuyu.dong@hpi.de

Patrick Lederer
ILLC, University of Amsterdam
Amsterdam, The Netherlands
p.lederer@uva.nl

Alexander Schlenga
Technical University of Munich
Munich, Germany
alexander.schlenga@tum.de

ABSTRACT

A central problem in multiagent systems is the fair assignment of objects to agents. In this paper, we initiate the analysis of classic majoritarian social choice functions in assignment. Exploiting the special structure of the assignment domain, we show a number of surprising results with no counterparts in general social choice. In particular, we establish a near one-to-one correspondence between preference profiles and majority graphs. This correspondence implies that key properties of assignments—such as Pareto-optimality, least unpopularity, and mixed popularity—can be determined solely by the associated majority graph. We further show that all Pareto-optimal assignments are semi-popular and belong to the top cycle. Elements of the top cycle can thus easily be found via serial dictatorships. Our main result is a complete characterization of the top cycle, which implies the top cycle can only consist of one, two, all but two, all but one, or all assignments. By contrast, we find that the uncovered set contains only very few assignments.

KEYWORDS

House Allocation; Social Choice Theory; Top Cycle; Uncovered Set

ACM Reference Format:
Felix Brandt, Haoyuan Chen, Chris Dong, Patrick Lederer, and Alexander Schlenga. 2026. Majoritarian Assignment Rules. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/QERE1891

1 INTRODUCTION

Assigning objects to individual agents is a fundamental problem that has received considerable attention by computer scientists as well as economists [e.g., 8, 18, 36, 44]. The problem is known as the *assignment problem*, the *house allocation problem*, or *two-sided matching with one-sided preferences*. In its simplest form, there are n agents, n houses, and each house needs to be allocated to exactly one agent based on the strict preferences of each agent over the houses. Applications are diverse and include assigning jobs to applicants, apartments to tenants, and offices to employees.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/QERE1891

A natural way to compare two assignments μ and λ in such settings is to check whether a majority of agents prefer the house they receive under μ to the one they receive under λ. This idea leads to the notion of *popular* assignments, which are assignments for which there is no other assignment that is preferred by a strict majority of the agents [27]. However, as Gärdenfors [27] pointed out, such popular assignments may not exist, as the majority relation over assignments can be cyclic. This observation has led to an extensive stream of follow-up work [e.g., 2, 3, 10, 14, 19, 20, 38], which aims to understand under which conditions popular assignments exist and which assignment should be selected in their absence. For example, the latter question has led to the definition of concepts such as *least unpopular* or *mixed popular* assignments [32, 37].

Curiously, this line of work closely resembles classic ideas in *social choice*. In this setting, agents report preferences over some alternatives and a social choice function (SCF) returns the "best" alternative based on these preferences. The observation that popular assignments do not always exist is analogous to the *Condorcet paradox* in social choice, which states that (weak) Condorcet winners need not exist. Even more, these two insights are logically related as house allocation can be viewed as a special case of social choice. When letting the set of alternatives be the set of all possible allocations of houses to agents and postulating that agents are indifferent between all allocations in which they receive the same house, house allocation is reduced to a restricted domain of general social choice. Under this reduction, a popular assignment is merely a weak Condorcet winner, so the absence of a popular assignment shows the Condorcet paradox for a restricted domain of preferences.[1] Moreover, several assignment concepts based on popularity have counterparts in social choice: for example, the notion of least popularity corresponds to the *maximin* voting rule and mixed popularity to *maximal lotteries*.

Motivated by this relationship, we aim to transfer further ideas from social choice, namely *majoritarian SCFs*, to the assignment domain. Just like popularity, such majoritarian SCFs only depend on the majority relation. Typical examples are Copeland's rule, the top cycle, the uncovered set, and the bipartisan set [e.g., 9, 35]. The definitions of these functions are equally natural in assignment as they are in social choice. The top cycle, for example, returns all assignments that are maximal elements of the transitive closure

[1] Majority cycles are even more frequent in assignment than they are in social choice, stressing the importance of which assignment to choose in the absence of popular ones. For example, this can be seen by comparing the proportions of profiles that admit Condorcet winners and popular assignments, respectively [2, 28, Tables 1 and 4.2].

of the majority relation. Similarly, the known set-theoretic relationships between these SCFs also hold in the assignment domain. Computational properties, on the other hand, do *not* carry over from social choice to assignment because, when viewing house allocation as a subdomain of social choice, the number of alternatives is exponential and the agent's preferences are concisely represented as each agent merely cares about her own house. This has serious algorithmic repercussions, and the computational complexity of even the simplest concepts needs to be reevaluated. For instance, identifying weak Condorcet winners is straightforward in social choice, but finding a popular assignment already requires clever algorithmic techniques [2].

Contribution

In this paper, we initiate the systematic study of majoritarian SCFs in the context of assignment. To this end, we first characterize all sets of preference profiles that admit the same majority graph by providing an efficient algorithm that reconstructs all profiles that induce a given majority graph. In sharp contrast to the social choice domain, it turns out that almost all majority graphs are induced by a *single* preference profile. As a consequence, the rules that return all Pareto-optimal assignments, all least unpopular assignments, and all mixed popular assignments, respectively, are majoritarian, even though these concepts are not defined in terms of the majority relation. Moreover, we show that all Pareto-optimal assignments are contained in the top cycle, which means that elements of the top cycle can be found via serial dictatorships. Further, Pareto-optimal assignments have non-negative Copeland score and are thus semi-popular. None of these results holds in the general social choice domain. We also prove that the rule returning all rank-maximal assignments is not majoritarian.

Our main result is a complete characterization of the top cycle in assignment when there are $n \geq 5$ agents and houses. This characterization shows that the cardinality of the top cycle may only take one of five values (1, 2, $n! - 2$, $n! - 1$, or $n!$) and leads to a simple sublinear-time algorithm that returns a concise representation of all assignments in the top cycle. Once again, this characterization has no analog in the social choice domain, where the top cycle can consist of any number of alternatives.

Lastly, we investigate the three most common variants of the uncovered set, all of which refine the top cycle. We compute the number of assignments contained in the uncovered set for $n = 5$ (by exhaustive enumeration) and for $n = 6$ and $n = 7$ (by sampling preference profiles). Somewhat surprisingly, in all these settings, most profiles only admit *two* uncovered assignments, suggesting that the uncovered set is much more discriminating in assignment than it is in social choice [cf. 16, 23].

Related Work

The study of matching under preferences was initiated by Gale and Shapley [26]. In their model (nowadays often referred to as marriage markets), there are two types of agents who have strict preferences over each other. Gale and Shapley showed that a so-called stable matching always exists and can be found by a simple, efficient algorithm. By contrast, Gärdenfors [27] showed that stable matchings are no longer guaranteed to exist when agents may be indifferent between other agents. As a remedy, he proposed the notion of popular matchings, i.e., matchings such that there is no other matching that a majority of the agents prefer. When individual preferences are strict, popularity is weaker than stability. However, popular matchings may not exist for weak preferences, even if agents of one type are completely indifferent between all agents of the other type and all other agents have strict preferences. This variant, which goes back to Gale [25], is known as assignment, house allocation, or two-sided matching with one-sided preferences as fully indifferent agents can be seen as objects. Abraham et al. [2] provided an efficient algorithm for finding a popular assignment or returning that none exists.

The definition of popularity has been relaxed in various ways to restore existence. McCutchen [37] proposed the unpopularity margin as a qualitative relaxation of popularity. A least unpopular matching in this sense always exists, but is NP-hard to compute. Kavitha et al. [32] introduced mixed popular assignments, whose existence is guaranteed by the minimax theorem, and provided an efficient algorithm for computing them. Kavitha and Vaish [33] transferred the notion of Copeland winners from social choice to so-called roommate markets, which generalize the assignment setting. They showed that computing a Copeland winner is NP-hard. It is open whether this hardness result also holds in assignment. Kavitha and Vaish also give a randomized FPTAS for identifying semi-popular matchings, another relaxation of popularity, in roommate markets. Semi-popular matchings are matchings that lose at most half of their majority comparisons [31]. Interestingly, one of our results, Proposition 1, entails that semi-popular *assignments* can be found efficiently via serial dictatorships in assignment.

Lastly, our work is related to the problem of voting on combinatorial domains [see, e.g., 34]. Two recent papers in this line of research that are close to ours are due to Brill et al. [17] and Boehmer and Dierking [4], who study matching through the lens of voting theory.

2 PRELIMINARIES

Let $n \in \mathbb{N}$ be given. We denote by $N = \{1, \dots, n\}$ a set of *agents* and by $H = \{a, b, c, \dots\}$ a set of n *houses* (or distinct indivisible objects in general). Each agent $x \in N$ reports a *preference relation* $\succ_x$, which is formally a linear order over H. Intuitively $a \succ_x b$ means that agent x prefers house a to house b. Note that we require each agent to rank all houses without indifference and that there are as many houses as agents. A *preference profile* $P = (\succ_1, \dots, \succ_n)$ is the collection of the preference relations of all agents.

Given a preference profile, our goal is to assign one house to each agent. To formalize this, we define *assignments* as bijective functions mapping agents in N to houses in H. Thus, $\mu(x)$ is the house given to agent x under assignment μ. We define by M the set of all possible assignments from N to H. Further, we write an assignment μ in which agents $1, 2, 3, \dots$ obtain houses $a, b, c, \dots$, respectively, as $\mu = (a, b, c, \dots)$. Throughout the paper, we assume that agents compare assignments only based on the houses they receive: an agent $x \in N$ (weakly) prefers an assignment μ to another assignment λ, denoted by $\mu \succeq_x \lambda$, if $\mu(x) \succ_x \lambda(x)$ or $\mu(x) = \lambda(x)$. Moreover, an agent $x \in N$ strictly prefers an assignment μ to an assignment λ, written as $\mu \succ_x \lambda$, if $\mu(x) \succ_x \lambda(x)$ and $\mu(x) \neq \lambda(x)$.

An *assignment rule* F maps every preference profile P to a non-empty set of assignments $F(P)$. The idea is that an assignment rule returns a set of "good" assignments, from which a single assignment will eventually be picked.

Most of the rules considered in this paper are *symmetric*, which demands that all agents and all houses are treated equally, respectively. More formally, relabeling the agents in the preference profile should correspond to relabeling the agents in the returned assignments, and relabeling the houses in all agents' rankings should correspond to relabeling the houses in the returned assignments. When $n \geq 2$, no assignment rule that always returns a single assignment can be symmetric. Therefore, we consider set-valued rules. To nonetheless distinguish more discriminating from less discriminating rules, we say that a rule F is a *refinement* of a rule G and write $F \subseteq G$, if $F(P) \subseteq G(P)$ for all profiles P.

Lastly, we discuss a standard property of assignments called Pareto-optimality. Intuitively, this notion requires that there is no assignment that makes one agent strictly better off without making another one worse off. To formalize this idea, we say an assignment μ *Pareto-dominates* another assignment λ in a profile P if all agents weakly prefer μ to λ and this preference is strict for at least one agent, i.e., $\mu \succeq_x \lambda$ for all agents $x \in N$ and $\mu \succ_x \lambda$ for at least one agent $x \in N$. Further, an assignment is *Pareto-optimal* if it is not Pareto-dominated by any other assignment. The set of Pareto-optimal assignments in a profile P is denoted by $PO(P)$.

The set of Pareto-optimal assignments is closely connected to the family of *serial dictatorships*. Such serial dictatorships are defined by a priority order $\sigma = (x_1, \dots, x_n)$ over the agents, and agents simply pick their favorite house that has not been taken yet in the order given by σ. It is known that an assignment is Pareto-optimal if and only if it is returned by a serial dictatorship for the given profile [1]. We note that Pareto-optimality is much more restrictive in assignment than in social choice: there are *always* Pareto-dominated assignments unless all agents have the same preferences. More generally, for every pair of agents x, y and every pair of houses a, b such that $a \succ_x b$ and $b \succ_y a$, there are $(n-2)!$ Pareto-dominated assignments where x obtains b and y obtains a.

3 THE STRUCTURE OF MAJORITY GRAPHS

A fundamental way to compare two assignments to each other is to postulate that one assignment is socially preferred to another if a majority of the agents prefer the former to the latter. To this end, let $N_{\mu,\lambda} = \{x \in N \mid \mu \succeq_x \lambda\}$ denote the set of agents who weakly prefer μ to λ. An assignment μ *weakly majority dominates* assignment λ if at least as many agents prefer μ to λ than vice versa, i.e., $\mu \succsim \lambda$ if $|N_{\mu,\lambda}| \geq |N_{\lambda,\mu}|$. Similarly, an assignment μ *(strictly) majority dominates* another assignment λ if strictly more agents prefer μ to λ than vice versa, i.e., $\mu \succ \lambda$ if $|N_{\mu,\lambda}| > |N_{\lambda,\mu}|$.

This naturally leads to the analysis of *majority graphs*, which have been extensively studied in social choice theory [e.g., 9, 35]. Specifically, the majority graph of a profile P is the directed graph $G_P = (M, \{(\mu, \lambda) \in M^2 \colon \mu \succsim \lambda\})$, which has all possible assignments as its vertices and there is an edge from an assignment μ to another assignment λ if μ weakly majority dominates λ.

Example 1. Consider the following profile P with $N = \{1, 2, 3\}$ and $H = \{a, b, c\}$, and the corresponding majority graph. A black arrow from μ to λ indicates that μ strictly majority dominates λ, whereas a bidirectional gray edge indicates a majority tie.

$$P = \begin{array}{l} 1\colon\; a,\; b,\; c \\ 2\colon\; a,\; b,\; c \\ 3\colon\; a,\; b,\; c \end{array}$$

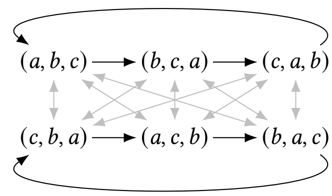

Note that we can obtain the same majority graph from other profiles, too. Specifically, if all three agents x rank $b \succ'_x c \succ'_x a$ in profile P', or all agents rank $c \succ''_x a \succ''_x b$ in profile P'', then $G_P = G_{P'} = G_{P''}$.

While every directed graph is induced by some preference profile in social choice [39], Brandt et al. [14] pointed out that this is not the case in assignment, where only a small fraction of majority graphs can actually be realized by preference profiles. Moreover, Brandt et al. gave an efficiently testable, necessary, and sufficient condition for two profiles yielding the same *weighted* majority graph, where each edge (λ, μ) of the majority graph is weighted by the margin $|N_{\lambda,\mu}| - |N_{\mu,\lambda}|$ of the majority comparison. In this section, we will generalize this result to unweighted majority graphs, showing that *almost all* majority graphs are induced by a *unique* profile.

To this end, we first recall some terminology by Brandt et al. [14]. Let P be a profile and let $(H_1, \dots, H_k)$ be an ordered partition of H, where we call each H_j a *component*. We say that $(H_j)_j$ is a *decomposition* of this profile, if all agents rank all houses in H_1 over all houses in H_2 and so on. Formally, for all $j < \ell \leq k$ and all $p \in H_j, q \in H_\ell$, and $x \in N$, it holds that $p \succ_x q$. Two profiles P, P' are called *rotation equivalent*, if the preferences on P and P' coincide within each component, and one ordering of the components is obtained by shifting the other. Formally, consider any decomposition $(H_1, \dots, H_k)$ of P. Then, for all $j \leq k$, $p, q \in H_j$, and $x \in N$, it has to hold that $p \succ_x q$ if and only if $p \succ'_x q$, and there exists some $r < k$ such that $(H_{1+r}, \dots, H_{k+r})$ is a decomposition of P' (where we set $H_{j+r} \coloneqq H_{j+r-k}$ if $j + r > k$).

Example 2. To illustrate rotation equivalence, consider the following profiles $\hat{P}, \overline{P}$ with $N = \{1, 2, 3, 4\}$, $H = \{a, b, c, d\}$, and decomposition $(H_1 = \{a\}, H_2 = \{b\}, H_3 = \{c, d\})$.

$$\hat{P} = \begin{array}{l} 1\colon\; a,\; b,\; c,\; d \\ 2\colon\; a,\; b,\; c,\; d \\ 3\colon\; a,\; b,\; d,\; c \\ 4\colon\; a,\; b,\; d,\; c \end{array} \qquad \overline{P} = \begin{array}{l} 1\colon\; a,\; b,\; c,\; d \\ 2\colon\; a,\; b,\; d,\; c \\ 3\colon\; a,\; b,\; d,\; c \\ 4\colon\; a,\; b,\; d,\; c \end{array}$$

The profile $\overline{P}$ is *not* rotation equivalent to $\hat{P}$, as $c \mathrel{\hat{\succ}_2} d$ and $d \mathrel{\bar{\succ}_2} c$ even though c and d belong to the same component H_3. Moreover, rotation equivalence can also be violated when, within each component, the preferences of the agents are coherent. For this, consider the profile P from Example 1 and let P''' be the profile, where all agents x report $a \succ'''_x c \succ'''_x b$. Then, P''' is not rotation equivalent to P, as the decomposition $(\{a\}, \{b\}, \{c\})$ cannot be rotated to $(\{a\}, \{c\}, \{b\})$. However, $(\{b\}, \{c\}, \{a\})$ and $(\{c\}, \{a\}, \{b\})$ are valid rotations. Hence, the profiles P' and P'' described in Example 1 are rotation equivalent to P.

Brandt et al. [14] showed that rotation equivalence characterizes the profiles that induce the same *weighted* majority graph. We are able to strengthen this result by showing that the margins are not required: rotation equivalence, in fact, characterizes the profiles inducing the same (unweighted) majority graph! As a consequence, given any assignment-induced majority graph, we can reconstruct the margins of all majority edges. Moreover, let a house "Pareto-dominate" another house if all agents rank the former above the latter. Whenever there are no Pareto-dominated houses in a profile P, our result implies that this profile has a unique majority graph G_P. Even in the presence of Pareto-dominated houses, we can deduce preferences of all agents except for the direction of some Pareto-dominations. The full proof of the following result is deferred to the full version [11, Appendix A].

Theorem 1. *Two profiles induce the same majority graph if and only if they are rotation equivalent.*

Proof sketch. It is easy to verify that rotation equivalent profiles indeed induce the same majority graph, so we focus on the remaining implication. Let G_{P^*} be a majority graph that is induced by some profile P^*. Our goal is to find all profiles P such that $G_P = G_{P^*}$. As a first step, we consider a pair of houses p, q. We iterate over pairs of agents x, y and instantiate an assignment μ in which x obtains p and y obtains q. We compare μ to the assignment λ in which x and y swap houses. If the two assignments create a majority tie, then x and y have identical preferences over p, q. However, if, e.g., μ is strictly majority-preferred to λ, then this means that $p \succ_x q$ and $q \succ_y p$ for any profile P with $G_P = G_{P^*}$. In other words, for each pair of distinct houses p and q, we can determine whether one Pareto-dominates the other (without knowing which one) for all profiles P with $G_P = G_{P^*}$. If this is not the case, then we can determine for each agent whether she ranks p over q in all such profiles P or vice versa.

Next, we instantiate a graph with the houses being the nodes. An edge between two houses p and q is added whenever not all agents prefer p to q or vice versa. Based on the insights from the previous paragraph, we can determine for each agent which of the houses they prefer more. This graph partitions the set of houses into connected components $H_1, \dots, H_k$. We then show that within each component H_i, we can determine the relative ordering between all pairs of houses by querying appropriate majority comparisons, and that each agent ranks the houses of each component contiguously. Finally, we re-order the components and prove that P can be decomposed as $(H_1, \dots, H_k)$ or a rotation $(H_{1+r}, \dots, H_{k+r})$ thereof. □

Remark 1. The proof of Theorem 1 yields an efficient algorithm for reconstructing all profiles inducing a given majority graph in time polynomial in n. Moreover, the majority graph uniquely determines the majority margins, which can also be deduced algorithmically. By contrast, to verify whether a given directed graph is the majority graph of a profile, one needs to check all $\sim(n!)^2$ edges of the graph.

Remark 2. Unless n is small, only very few profiles admit a non-trivial decomposition, implying that they can be fully reconstructed from their majority graph. As a matter of fact, almost all majority graphs are induced by a *single* preference profile. Calculations by Brandt et al. [14] demonstrate that more than 99% of all majority graphs are induced by a single profile as soon as $n \geq 4$.

4 MAJORITARIAN ASSIGNMENT RULES

The concept of majority graphs has given rise to numerous influential solution concepts in social choice theory, such as Condorcet winners, Copeland's rule, the top cycle, and the uncovered set [see, e.g., 9, 35]. In particular, all of these concepts are majoritarian, i.e., they can be computed solely based on the majority graph of a profile. As a consequence, the definitions of these concepts directly carry over to the assignment domain while preserving their natural appeal. Weak Condorcet winners, for example, are known as popular assignments in house allocation. An assignment μ is *popular* if $\mu \succsim \lambda$ for all $\lambda \in M$. Example 1 shows that popular assignments need not exist. In the following, we investigate *majoritarian* assignment rules, i.e., assignment rules that only depend on the majority graph. Formally, an assignment rule F is majoritarian if $F(P) = F(P')$ for all profiles P and P' with $G_P = G_{P'}$.

While our main focus is the study of established majoritarian voting rules in the context of assignment, Theorem 1 implies that several well-known assignment concepts are actually majoritarian. Specifically, this result entails that an assignment rule is majoritarian if and only if it is invariant with respect to rotation equivalence. We use this fact to prove that Pareto-optimality, least unpopularity, and mixed popularity are majoritarian. By *least unpopularity*, we denote the rule that returns all assignments minimizing the margin of their worst majority defeat. By *mixed popularity*, we denote the rule that returns all assignments which are part of the support of some mixed popular matching. Formal definitions of these concepts can, for example, be found in the papers by McCutchen [37], Kavitha et al. [32], and Brandt and Bullinger [10].

Corollary 1. *PO, least unpopularity, and mixed popularity are majoritarian.*

Proof. Let P and P' be two rotation equivalent profiles w.r.t. some decomposition $(H_1, \dots, H_k)$ and a shift by $r \in \{1, \dots, k-1\}$.

PO: Let $\mu \in PO(P)$ be a Pareto-optimal assignment. By a characterization due to Abdulkadiroğlu and Sönmez [1], there is an order over the agents $\sigma = (x_1, \dots, x_n)$ such that μ is chosen by the serial dictatorship SD_σ induced by σ, i.e., $\mu = SD_\sigma(P)$. Note that the agents choose houses from $H_1, \dots, H_k$ in this order because, for all $i, j \in \{1, \dots, k\}$ with $i < j$, it holds that every agent prefers every house $h \in H_i$ to every house $h' \in H_j$. Next, we partition the agents $x \in N$ into the sets $N_i = \{x \in N : \mu(x) \in H_i\}$ for $i \in \{1, \dots, k\}$. Let σ' denote the order of agents such that *(i)* for all $i, j \in \{1, \dots, k\}$ with $i < j$, all agents in N_{i+r} are ranked before all agents in N_{j+r} and *(ii)* within each set N_i, the agents are ordered the same as in σ. Under the serial dictatorship $SD_{\sigma'}$ induced by this sequence, the agents from N_{1+r} first get to choose their houses. Since P' is achieved by rotating P with a shift of r, these agents obtain precisely the same houses from H_{1+r} as in μ. Inductively, the agents in N_{j+r} obtain precisely the houses from H_{j+r} under this shifted picking sequence in the profile P', and the obtained assignment is hence μ. This proves that $\mu \in PO(P')$. Reversing the roles of P and P', we obtain that $PO(P) = PO(P')$.

Least unpopularity and mixed popularity: The proof of Theorem 1 shows that the margins of all majority comparisons can be inferred from the majority graph. Since least popularity and mixed popularity only depend on these margins, they are majoritarian. □

By contrast, we show next via an example that the rule that returns all *rank-maximal assignments* [30] is not majoritarian. To introduce rank-maximality, we define the rank of a house p in a preference relation $\succ$ by $r(\succ, p) = 1 + |\{q \in H \mid q \succ p\}|$. Further, the rank vector of an assignment μ for a profile P contains the ranks $r(\succ_x, \mu(x))$ of each agent $x \in N$ for her assigned house in increasing order. Then, an assignment μ is rank-maximal if its rank vector is lexicographically optimal, i.e., the assignment maximizes the number of agents who obtain their favorite house, subject to this condition it maximizes the number of agents who obtain their second-ranked house, and so on.

Example 3. Consider the following two profiles P and P'. They are rotation equivalent by the decomposition $(\{a, b, c\}, \{d, e, f\})$ and thus induce the same majority graph.

$$P = \begin{array}{ll} 1: & a,\ b,\ c,\ d,\ e,\ f \\ 2: & a,\ c,\ b,\ d,\ f,\ e \\ 3: & b,\ a,\ c,\ e,\ d,\ f \\ 4: & a,\ b,\ c,\ d,\ e,\ f \\ 5: & a,\ b,\ c,\ d,\ e,\ f \\ 6: & a,\ b,\ c,\ d,\ e,\ f \end{array} \qquad P' = \begin{array}{ll} 1: & d,\ e,\ f,\ a,\ b,\ c \\ 2: & d,\ f,\ e,\ a,\ c,\ b \\ 3: & e,\ d,\ f,\ b,\ a,\ c \\ 4: & d,\ e,\ f,\ a,\ b,\ c \\ 5: & d,\ e,\ f,\ a,\ b,\ c \\ 6: & d,\ e,\ f,\ a,\ b,\ c \end{array}$$

The assignment $\mu = (a, c, b, d, e, f)$ marked in blue is rank maximal in P, but not in P', as the assignment $\lambda = (d, f, e, a, b, c)$ in red assigns two agents their top choice. To see that μ indeed is rank-maximal in P, note that any assignment can give at most two agents their top choices. Agent 3 has to obtain b, and since 1, 4, 5, 6 all have the same preferences, we can assign a to agent 1. Among the remaining agents, only 2 can still obtain her second-favorite house, c. Among 4, 5, 6, it then does not matter how we assign d, e, f for rank-maximality.

We conclude this section by proving another surprising relationship between two majoritarian assignment rules. Specifically, we show that all Pareto-optimal assignments are semi-popular. Semi-popularity is a weakening of popularity, which requires that an assignment is majority preferred to at least half of all assignments. More formally, an assignment μ is *semi-popular* in a profile P if $|\{\lambda \in M \mid \mu \succsim \lambda\}| \geq \frac{|M|}{2}$ [33]. By $SP(P)$, we denote the set of all semi-popular assignments in a profile P.

Proposition 1. *$PO \subseteq SP$.*

Proof. For our proof, we first introduce permutations on assignments. Given a permutation π on N, define $\pi' : (N \to H) \to (N \to H)$ such that for any assignment $\mu : N \to H$ and agent $x \in N$, we have $\pi'(\mu)(x) := \mu(\pi(x))$. Intuitively, the assignment $\mu' = \pi'(\mu)$ is obtained from μ by assigning to agent x the house that is given to agent $\pi(x)$ in μ. For the sake of simplicity, we slightly abuse notation and refer to π' as π too.

Now, fix an arbitrary profile P and an assignment $\mu \in PO(P)$. We consider an arbitrary permutation π and show that $\pi(\mu) \succ \mu$ implies that $\mu \succ \pi^{-1}(\mu)$. For simplicity, we name $\pi(\mu) =: \eta$ and $\pi^{-1}(\mu) =: \lambda$. Let y denote an arbitrary agent who strictly prefers η to μ, i.e., $\eta(y) \succ_y \mu(y)$. Further, let $x = \mu^{-1}(\eta(y))$ be the agent who gets $\eta(y)$ in μ. Note that $x \neq y$. Since μ is Pareto-optimal, it cannot be that $\mu(y) \succ_x \mu(x)$, as otherwise swapping the houses of x and y would be a Pareto-improvement over μ. Therefore, $\mu(x) \succ_x \mu(y) = \lambda(x)$ and $\mu \succ_x \lambda$. Since y was chosen arbitrarily, we see that for every agent strictly preferring η to μ, we have one other strictly preferring μ to λ. Moreover, if $\mu(x) = \eta(x)$, then $\pi(x) = x$, which implies also that $\mu(x) = \lambda(x)$. Hence, if a majority of agents prefer η to μ, a majority of agents prefer μ to λ. We lastly note that every assignment η can be obtained by permuting μ, i.e., there is some permutation π such that $\eta = \pi(\mu)$. Hence, it follows that $|\{\lambda \in M \mid \mu \succsim \lambda\}| \geq \frac{|M|}{2}$, so μ is semi-popular. □

4.1 The Top Cycle

We next turn to the top cycle, one of the most prominent majoritarian rules in the social choice domain [e.g., 5, 15, 29, 41, 43, 45]. The underlying idea is very natural: popular assignments do not always exist because the majority relation $\succsim$ fails to be transitive (see Example 1). Instead, one can consider $\succsim^*$, the *transitive closure* of $\succsim$, and simply return the maximal elements according to this relation. Formally, $TC(P) = \{\mu \in M : \forall \lambda \in M : \mu \succsim^* \lambda\}$.[2,3] Or, in other words, the top cycle returns all assignments that reach every other assignment on some path in the majority graph.

As a first step towards understanding the top cycle in the assignment domain, we prove that it always contains all Pareto-optimal assignments. This is not true in the social choice domain [see, e.g., 35, Theorem 10.2.3]. The proof of the following proposition as well as Theorem 2 can be found in the full version [11, Appendix B].

Proposition 2. *$PO \subseteq TC$.*

Using Proposition 2 as a stepping stone, we obtain a much stronger structural result about majority graphs in assignment: when $n \geq 5$, the top cycle can only contain one, two, all but one, all but two, or all assignments.

Theorem 2. *Let P be any profile with $n \geq 5$ agents and houses. Then, $|TC(P)| \in \{1, 2, n! - 2, n! - 1, n!\}$. More precisely, we have*

(i) $|TC(P)| = 1$ if all agents have distinct top choices,[4]
(ii) $|TC(P)| = 2$ if all but two agents have distinct top choices. Further, these two also share the same second choice, which is not top-ranked by any other agent either,
(iii) $|TC(P)| = n! - 2$ if (i) and (ii) do not hold, and all but two agents have distinct bottom choices. Further, these two also share the same second-to-bottom choice, which is not last-ranked by any other agent either,
(iv) $|TC(P)| = n! - 1$ if (i) and (ii) do not hold, and all agents have distinct bottom choices, and
(v) $|TC(P)| = n!$ if none of the above cases holds.

Proof sketch. First, the cases *(i)* and *(ii)* follow relatively easily from Proposition 2: we get under the corresponding assumptions that there are 1 or 2 Pareto-optimal assignments, and it is easy to show that these are the only ones in the top cycle. Moreover, if the corresponding assumptions are not true, we can show that there are more than 1 (resp. 2) Pareto-optimal assignments.

[2] One can also consider $\succ^*$, the transitive closure of the *strict* part of the majority relation and return its maximal elements. The resulting SCF is known as the Schwartz set or GOCHA [41]
[3] The top cycle is unrelated to the *Top Trading Cycle (TTC)* by Shapley and Scarf [42], an assignment algorithm for settings with initial endowments.
[4] This case corresponds to the profiles admitting a strongly popular assignment, which strictly majority dominate every other assignment. These are known as (strict) Condorcet winners in social choice.

For cases *(iii)* through *(v)*, we first prove via a computer-aided approach that our result holds when there are $n = 5$ agents and houses. Specifically, we let the computer enumerate all profiles for $n = 5$ (up to symmetries) and verify that our theorem is true. Based on this insight, we then tackle the remaining cases when $n \geq 6$. To this end, we define the bottom cycle $BC(P)$ as the set of assignments that can be reached from every other assignment via a path in the majority graph. Further, let $P_{N',H'}$ denote the restriction of our input profile P to a set of agents N' and set of houses H' (with $|N'| = |H'|$), and, for every assignment μ, let $\mu_{N',H'}$ be the assignment from N' to H' induced by μ. Based on our five-agent case, we then prove that there is a path from λ to μ in the majority relation if the sets $N' = \{x \in N : \mu(x) \neq \lambda(x)\}$ and $H' = \mu(N')$ satisfy that $|N'| = |H'| = 5$, $\lambda_{N',H'} \notin BC(P_{N',H'})$ or $|BC(P_{N',H'})| > 2$, and $\mu_{N',H'} \notin TC(P_{N',H'})$ or $|TC(P_{N',H'})| > 2$.

Next, we show that, under cases *(iii)*, *(iv)*, and *(v)*, there is an agent x^* such that, if x^* get his most preferred house p in an assignment μ, then μ belongs to the top cycle. To prove this claim, we consider a case distinction regarding the shape of P. In particular, since cases *(i)* and *(ii)* do not apply, we know that either (1) there is a house p that is top-ranked by at least three agents, (2) there are two houses p and q that are both top-ranked by two agents, or (3) there are two houses p and q such that two agents top-rank p, one of them second-ranks q, and a third agent top-ranks q.

For example, to prove our claim for case (1), let x, y, z denote three agents that top-rank p and let μ denote an assignment derived from a serial dictatorship where x picks first. By this definition, $\mu(x) = p$ and $\mu \in TC(P)$ because μ is Pareto-optimal. Next, let λ denote an assignment derived from μ by swapping the houses of two arbitrary agents $u, v \in N \setminus \{x\}$. We claim that $\lambda \in TC(P)$. To see this, we choose a set $N' \supseteq \{x, y, z, u, v\}$ with $|N'| = 5$ and let $H' = \lambda(N)$. In $P_{N',H'}$, the agents x, y, z still top-rank p, so we infer from cases (i) and (ii) that $|TC(P_{N',H'})| > 2$. Further, since $\lambda(x) = \mu(x) = p$, $\lambda_{N',H'} \notin BC(P_{N',H'})$ unless $|BC(P_{N',H'})| > 2$. Hence, the insights of the previous paragraph show that $\lambda \succsim^* \mu$. By repeatedly applying this construction, it follows that every assignment λ with $\lambda(x) = p$ is in the top cycle. The other cases rely on similar ideas.

As the next step, we prove that $TC(P)$ contains all assignments except for so-called Pareto-pessimal ones, which are those that do not Pareto-dominate any other assignment. To prove this claim, we focus on an assignment μ that Pareto-dominates another assignment λ. Hence, for all agents x with $\mu(x) \neq \lambda(x)$, it holds that $\mu(x) \succ_x \lambda(x)$. Further, let x^* denote the agent such that, if x^* gets his favorite house p in an assignment, the assignment is in the top cycle. We show that we can modify λ such that agent x^* is assigned his favorite house p while maintaining that μ majority dominates the resulting assignment η. Since $\eta \in TC(P)$ by our previous discussion, this proves that $\mu \in TC(P)$, too.

Finally, we observe that the top cycle of a profile is the bottom cycle of the profile where we reversed all agents' preference relations. This implies that all Pareto-pessimal assignments are contained in the bottom cycle and that the bottom cycle contains all assignments that are not Pareto-optimal if $|BC(P)| > 2$. For cases *(iii)* and *(iv)*, our theorem then follows by showing that there are precisely two or one Pareto-pessimal assignments. By contrast, for case *(v)*, we show that either all assignments are Pareto-optimal and thus in TC, or there is an assignment that is not Pareto-optimal but in the top cycle. Since this assignment is also in the bottom cycle in this case, $TC(P) = BC(P)$, which implies that $TC(P) = M$. ◻

Theorem 2 shows that deterministic assignments are highly unstable with respect to majority deviations. Indeed, given a starting assignment and a target assignment, one can almost always convince the agents to transition from one to the other by presenting intermediate assignments that are weakly majority preferred. To illustrate this point, consider the following poor assignment that reaches every other assignment via some majority path.

Example 4. In the subsequent profile P, the assignment μ marked in red is obviously not desirable. It fails to be Pareto-optimal, and many agents even receive their least-preferred house. Nevertheless, it is contained in the top cycle. A path of dominations via which μ reaches a serial dictatorship (and by virtue of Proposition 2 the entire top cycle) is given in the full version [11, Appendix C]. The rough idea is to repeatedly reassign to some agents slightly worse houses while improving other agents' assignments significantly. For example, we can make agents 4, 5 worse by assigning c, a to them, respectively. However, this frees houses d, e which we assign to agents 2 and 7, thereby making these agents significantly happier.

$$P = \begin{array}{l} 1: \; f, \; b, \; d, \; e, \; c, \; a, \; \color{red}{g} \\ 2: \; d, \; f, \; g, \; a, \; b, \; e, \; \color{red}{c} \\ 3: \; d, \; a, \; c, \; g, \; e, \; b, \; \color{red}{f} \\ 4: \; a, \; d, \; b, \; f, \; g, \; \color{red}{e}, \; c \\ 5: \; c, \; g, \; e, \; b, \; f, \; \color{red}{d}, \; a \\ 6: \; f, \; a, \; e, \; d, \; g, \; c, \; \color{red}{b} \\ 7: \; c, \; d, \; e, \; b, \; g, \; f, \; \color{red}{a} \end{array}$$

Remark 3. For completeness, we also consider the cases $n < 5$ for Theorem 2. Clearly, there exists only one assignment for $n = 1$, and two assignments for $n = 2$. For $n = 3$, we have found via a computer-aided approach that the top cycle has size either 1, 2, $n! - 2 = 4$, or $n! = 6$. However, in contrast to case *(iv)* of Theorem 2, the top cycle may contain $n! - 2$ assignments even though all agents have pairwise distinct bottom choices. This happens, for example in the following profile, where $\mu = (b, c, a)$ (in red) and $\lambda = (c, b, a)$ do not belong to the top cycle.

$$P = \begin{array}{l} 1: \; a, \; \color{red}{b}, \; c \\ 2: \; a, \; \color{red}{c}, \; b \\ 3: \; c, \; b, \; \color{red}{a} \end{array}$$

Lastly, for $n = 4$, we found via our computer-aided approach that the top cycle can, in addition to the five sizes described in Theorem 2, also have a size of $n! - 3$. Up to symmetries, this happens precisely in the following two profiles P and P', where the assignments $\mu = (c, d, b, a)$ (marked in red), $\lambda = (d, c, a, b)$, and $\eta = (d, c, b, a)$ are respectively not contained in the top cycle.

$$P = \begin{array}{l} 1: \; a, \; b, \; \color{red}{c}, \; d \\ 2: \; a, \; b, \; \color{red}{d}, \; c \\ 3: \; c, \; d, \; a, \; \color{red}{b} \\ 4: \; c, \; d, \; b, \; \color{red}{a} \end{array} \qquad P' = \begin{array}{l} 1: \; a, \; b, \; \color{red}{c}, \; d \\ 2: \; a, \; b, \; \color{red}{d}, \; c \\ 3: \; d, \; c, \; a, \; \color{red}{b} \\ 4: \; d, \; c, \; b, \; \color{red}{a} \end{array}$$

Remark 4. Theorem 2 stands in stark contrast to classic social choice, where the top cycle may have any number of elements, even when there are at most three agents. Further, in social choice theory, TC can be computed in linear time in the size of the profile [9, 13].

Theorem 2 implies that in assignment, computing and returning a concise representation of the (possibly exponentially large) top cycle is possible in *sub-linear* time.

4.2 Uncovered Sets

As our final contribution, we turn to another technique addressing the non-transitivity of the majority relation: uncovered sets. These sets are based on covering relations, which are natural *transitive subrelations* of the majority relation. Just as in the definition of the top cycle, we can take the maximal elements for each of these relations, defining an uncovered set that refines the top cycle. Just as the top cycle, uncovered sets have been extensively studied in social choice theory [see, e.g., 6, 7, 12, 21].

The presence of majority ties in the assignment domain allows for multiple definitions of covering relations and uncovered sets, and we will subsequently define the three most common ones. Given a profile P, an assignment μ *Bordes covers* another assignment λ, if $\mu \succ \lambda$, and for every third assignment η, we have that $\lambda \succ \eta$ implies $\mu \succ \eta$. Similarly, μ *Gillies covers* λ, if $\mu \succ \lambda$, and for every η, we have that $\eta \succ \mu$ implies $\eta \succ \lambda$. Finally, μ *McKelvey covers* λ if it Bordes and Gillies covers it. Each of the three covering relations gives rise to a corresponding *uncovered set* (*UC*). It returns the maximal assignments of the covering relation, i.e., $UC(P) = \{\mu \in M\colon \text{no } \lambda \in M \text{ covers } \mu\}$. Whenever we refer to covering or *UC* without further specification, we mean McKelvey covering. All three uncovered sets can be characterized as assignments that reach all other assignments via some majority path of length at most 2. For Bordes, the first segment of any path of length 2 must be strict; for Gillies, the second segment must be strict; and for McKelvey, one of the two segments must be strict. This immediately implies that all uncovered sets are contained in the top cycle. Moreover, both the Bordes and the Gillies uncovered set are refinements of the McKelvey uncovered set.

From the general social choice setting, we know that $UC \subseteq PO$ [24]. This inherits to assignment, and we can easily prove that the inclusion is strict on this domain, too.

Example 5. In the following profile, the assignment $\mu = (c, a, b)$ in blue McKelvey-covers the assignment in red $\lambda = (a, b, c)$, even though λ is Pareto-optimal.

$$P = \begin{array}{rl} 1\colon & a,\ c,\ b \\ 2\colon & a,\ b,\ c \\ 3\colon & b,\ a,\ c \end{array}$$

Recall that all serial dictatorships are Pareto-optimal. Hence, this example illustrates that *PO* fails to distinguish between "good" picking sequences and "bad" ones in which the agents take away each other's favorite houses in unfortunate ways. This effect occurs for arbitrarily large numbers of agents. Thus, the set of uncovered assignments can be seen as a particularly attractive subset of the set of Pareto-optimal assignments.

To see how decisive *UC* is, we computed its choice sets and tracked the occurring cardinalities while iterating over all preference profiles up to symmetry for $n = 5$ exhaustively. The resulting graph is depicted in Figure 1. Further, we sampled profiles for $n = 7$ agents drawing each agent's preferences uniformly at random, depicted in Figure 2. It turns out that the Bordes-*UC* is almost

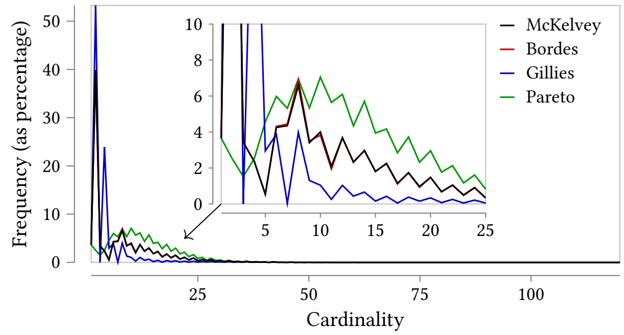

Figure 1: Size distributions of UCs for $n = 5$. The high peak is at size 2 for all of them. In total, there are 9,078,630 profiles (up to symmetries) and there are $5! = 120$ different assignments.

indistinguishable from the McKelvey-*UC*, while the Gillies-*UC* is, on average, the most discriminating one. This can be explained as follows: Under the Gillies-*UC*, for an assignment λ to not be covered despite $\mu \succ \lambda$, there needs to exist another assignment η such that $\lambda \succsim \eta \succ \mu$. However, for small numbers of agents, there are many profiles admitting popular assignments. If μ is such a popular assignment, there exists no η with $\eta \succ \mu$, and hence μ automatically Gillies-covers all λ with $\mu \succ \lambda$. Most notably, many profiles in both simulations admitted an uncovered set of size two. This finding suggests that *UC* is much more discriminative in assignment than in general social choice. However, can this already be explained by *PO* being more discriminative in assignment than in social choice?

To investigate how much *UC* differs from *PO*, we exhaustively studied the case when $n = 5$. For this, we utilize that the rules are symmetric with respect to permuting agents and houses. We therefore fix the preferences of agent 1, and further demand that the preferences of agents 2 through 5 are ordered lexicographically. This results in roughly 9 million profiles to be checked, which can be done within a few days on a computer. In Figure 3, we depict the percentage of profiles P for which the ratio $\frac{|UC(P)|}{|PO(P)|}$ is at most $x \in [0, 1]$, as a function of x. The results suggest that, indeed, *UC* is significantly more discriminative than *PO*. Hence, this consolidates that *UC* is an interesting refinement of *PO*, and it seems worthwhile to further investigate the properties of *UC*.

Remark 5. In social choice, *UC* can be computed in polynomial time via matrix multiplication [9, 13]. In assignment, the uncovered set can be exponentially large. Unless a structural result such as Theorem 2 also holds for *UC*, it seems unlikely that the uncovered set can be returned in polynomial time. Instead, the two interesting questions are *(i)* whether, given a profile, one can efficiently find an uncovered assignment, and *(ii)* whether, given a profile and an assignment, the assignment is uncovered. The computational complexity of both problems remains open.

Remark 6. Our experiments suggest that, at least for small n, all rank-maximal assignments are contained in the uncovered set. We have verified this through exhaustive search for all profiles with $n \leq 5$. If this were true in general, one could obtain an element of the uncovered set by computing a rank-maximal assignment,

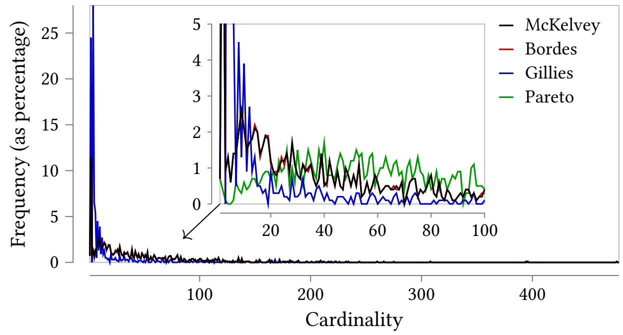

Figure 2: Size distributions of UCs for $n = 7$ in $1,000$ profiles sampled via the impartial culture model. The high peak is at size 2 for McKelvey and Bordes. Gillies-*UC* has an even higher peak at size 4.

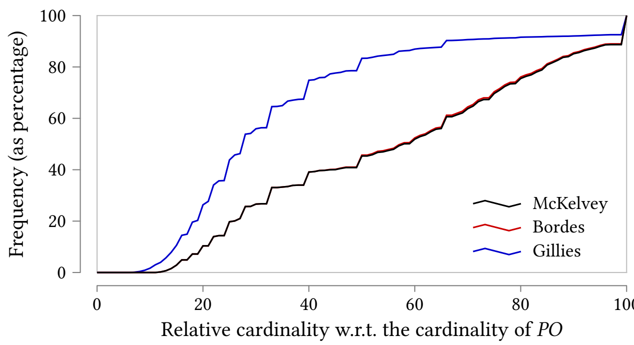

Figure 3: Cumulative distributions of the ratio of UC and PO sizes for $n = 5$. The plot shows the relative difference, i.e., the ratio. In total, there are $9,078,630$ profiles up to symmetry. We note that the plots for McKelvey and Bordes almost perfectly align.

which is possible in polynomial time. A set inclusion between the two rules would be interesting, as then *UC* would be a natural rule that is relatively decisive and yet contains all rank-maximal *and* all popular assignments.

Remark 7. Generous assignments [36, 40] are a dual version of rank maximal assignments: we compare assignments again based on their rank vector, but we now lexicographically optimize for the worst-off agents. While generous and rank-maximal assignments have very similar definitions, *UC* indicates that rank-maximal assignments may be preferable, as there are simple instances where no generous assignment is in the *UC*. For example, consider the following profile P.

$$P = \begin{array}{llllllll} 1: & \underline{c} & f & a & e & b & g & d \\ 2: & \underline{b} & c & g & e & a & d & f \\ 3: & g & f & a & \underline{d} & e & c & b \\ 4: & \underline{g} & b & e & c & a & f & d \\ 5: & \underline{e} & d & a & b & c & f & g \\ 6: & \underline{a} & b & d & g & f & e & c \\ 7: & \underline{f} & b & d & e & c & a & g \end{array}$$

The only two uncovered assignments for P are the underlined $\mu = (c, b, d, g, e, a, f)$ and the blue $\lambda = (c, b, g, d, e, a, f)$. However, neither of these is generous, as we can modify μ by giving agent 3 house a and agent 6 house d. This modified assignment μ' gives all agents houses within their top three, while both μ and λ give some agent her fourth-best house or worse. Hence, no generous assignment in this profile is uncovered.

5 CONCLUSION AND FUTURE WORK

In this paper, we initiate a systematic study of majoritarian assignment rules—set-valued assignment rules that rely solely on the pairwise majority relation. Prior work on majoritarian concepts in the context of assignment was restricted to popularity, corresponding to weak Condorcet winners in social choice theory.[5] However, just like weak Condorcet winners, popular assignments rarely exist. To circumvent this issue, social choice theory has developed a range of majoritarian functions that return sets of "good" alternatives in the absence of Condorcet winners. We have transferred two of the most prominent such functions—the top cycle and the uncovered set—to the subdomain of assignment. These rules are symmetric, treating all agents and houses equally, and they help to narrow down the set of acceptable assignments, from which a final selection (e.g., by randomization) can be made.

We proved a structural result about assignment-induced majority graphs, which, surprisingly, revealed that some well-known assignment rules are majoritarian. We then gave a complete and efficiently checkable characterization of the assignments contained in the top cycle. This characterization reveals that the top cycle not only contains all Pareto-optimal assignments (which does not hold in the more general social choice domain) but also some rather unattractive ones. The top cycle is too coarse to exclude these undesirable assignments. By contrast, the three variants of the uncovered set we studied are much more selective. In fact, each of them contains a subset of all Pareto-optimal assignments and thus offers a promising foundation for new, appealing assignment rules.

Our findings pave the way for the exploration of further appealing refinements of the McKelvey uncovered set, such as the minimal covering set and the bipartisan set (aka sign-essential set) [see, e.g., 9, 12, 22]. Both of these rules can be computed efficiently in social choice theory. Whether this is also true in assignment is wide open. Indeed, even seemingly simpler problems—such as finding an assignment in the uncovered set or the Copeland set, or deciding whether a given assignment belongs to any of these sets—remain unresolved. In particular, it would be interesting to investigate whether—in contrast to social choice—the support of mixed popular assignments is contained in the uncovered set.

Other avenues for future research include relaxations of the model that allow for different numbers of agents and houses, ties in the preferences, and pairwise matchings of agents. These generalizations would broaden the applicability of majoritarian assignment rules and deepen our understanding of their properties.

[5] Kavitha and Vaish [33] have studied semi-popularity and Copeland winners in the more general setting of roommate markets.

ACKNOWLEDGMENTS

This work was supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/11-2. Most of this research was carried out while Patrick Lederer was at UNSW Sydney, where he was funded by the NSF-CSIRO project on "Fair Sequential Collective Decision Making" (RG230833). Similarly, Chris Dong was at TU Munich, supported by the above grant of the Deutsche Forschungsgemeinschaft. We greatly appreciate the contributions of Martin Bullinger and Matthias Greger during our initial research meetings on this topic. Parts of this work were presented at the rump session of the PRAGMA Fest 2025 in Krakow.

REFERENCES

[1] 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, Robert W. Irving, Telikepalli Kavitha, and K. Mehlhorn. 2007. Popular matchings. *SIAM J. Comput.* 37, 4 (2007), 1030–1034.

[3] Haris Aziz, Felix Brandt, and Paul Stursberg. 2013. On Popular Random Assignments. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT) (Lecture Notes in Computer Science (LNCS), Vol. 8146)*. Springer-Verlag, 183–194.

[4] Niclas Boehmer and Jessica Dierking. 2025. *Single-Winner Voting on Matchings*. Technical Report. https://arxiv.org/pdf/2601.19653.

[5] Georges Bordes. 1976. Consistency, Rationality and Collective Choice. *Review of Economic Studies* 43, 3 (1976), 451–457.

[6] Georges Bordes. 1983. On the Possibility of Reasonable Consistent Majoritarian Choice: Some Positive Results. *Journal of Economic Theory* 31, 1 (1983), 122–132.

[7] Georges Bordes, Michel Le Breton, and M. Salles. 1992. Gillies and Miller's Subrelations of a Relation over an Infinite Set of Alternatives: General Results and Applications to Voting Games. *Mathematics of Operations Research* 17, 3 (1992), 509–518.

[8] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, J. Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 12.

[9] Felix Brandt, Markus Brill, and Paul Harrenstein. 2016. Tournament Solutions. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, J. Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 3.

[10] Felix Brandt and Martin Bullinger. 2022. Finding and Recognizing Popular Coalition Structures. *Journal of Artificial Intelligence Research* 74 (2022), 569–626.

[11] Felix Brandt, Haoyuan Chen, Chris Dong, Patrick Lederer, and Alexander Schlenga. 2026. *Majoritarian Assignment Rules*. Technical Report. https://arxiv.org/abs/2602.14816.

[12] Felix Brandt and Felix Fischer. 2008. Computing the Minimal Covering Set. *Mathematical Social Sciences* 56, 2 (2008), 254–268.

[13] Felix Brandt, Felix Fischer, and Paul Harrenstein. 2009. The Computational Complexity of Choice Sets. *Mathematical Logic Quarterly* 55, 4 (2009), 444–459. Special Issue on Computational Social Choice.

[14] Felix Brandt, Johannes Hofbauer, and Martin Suderland. 2017. Majority Graphs of Assignment Problems and Properties of Popular Random Assignments. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 335–343.

[15] Felix Brandt and Patrick Lederer. 2023. Characterizing the Top Cycle via Strategyproofness. *Theoretical Economics* 18, 2 (2023), 837–883.

[16] Felix Brandt and Hans Georg Seedig. 2016. On the Discriminative Power of Tournament Solutions. In *Selected Papers of the International Conference on Operations Research, OR2014*. Springer-Verlag, 53–58.

[17] Markus Brill, Niclas Boehmer, and Ulrike Schmidt-Kraepelin. 2025. Proportional representation in matching markets: selecting multiple matchings under dichotomous preferences. *Social Choice and Welfare* 64, 1–2 (2025), 179–220.

[18] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. 2006. Issues in Multiagent Resource Allocation. *Informatica* 30 (2006), 3–31.

[19] Ágnes Cseh. 2017. Popular Matchings. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 6.

[20] Ágnes Cseh, Chien-Chung Huang, and Telikepalli Kavitha. 2015. Popular matchings with two-sided preferences and one-sided ties. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP) (Lecture Notes in Computer Science (LNCS), Vol. 9134)*. Springer-Verlag, 367–379.

[21] John Duggan. 2013. Uncovered Sets. *Social Choice and Welfare* 41, 3 (2013), 489–535.

[22] Bhaskar Dutta and Jean-François Laslier. 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare* 16, 4 (1999), 513–532.

[23] Mark Fey. 2008. Choosing From a Large Tournament. *Social Choice and Welfare* 31, 2 (2008), 301–309.

[24] Peter C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.

[25] David Gale. 1960. *The Theory of Linear Economic Models*. McGraw-Hill.

[26] David Gale and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

[27] Peter Gärdenfors. 1975. Match Making: Assignments based on bilateral preferences. *Behavioral Science* 20, 3 (1975), 166–173.

[28] William V. Gehrlein and Peter C. Fishburn. 1979. Proportions of Profiles with a Majority Candidate. *Computers & Mathematics with Applications* 5, 2 (1979), 117–124.

[29] I. J. Good. 1971. A Note on Condorcet Sets. *Public Choice* 10, 1 (1971), 97–101.

[30] Robert W. Irving, Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E. Paluch. 2006. Rank-maximal matchings. *ACM Transactions on Algorithms* 2, 4 (2006), 602–610.

[31] Telikepalli Kavitha. 2020. Min-Cost Popular Matchings. In *Proceedings of the 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*. 25:1–25:17.

[32] Telikepalli Kavitha, Julián Mestre, and M. Nasre. 2011. Popular mixed matchings. *Theoretical Computer Science* 412, 24 (2011), 2679–2690.

[33] Telikepalli Kavitha and Rohit Vaish. 2023. Semi-Popular Matchings and Copeland Winners. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 957–965.

[34] Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, J. Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, Chapter 9.

[35] Jean-François Laslier. 1997. *Tournament Solutions and Majority Voting*. Springer-Verlag.

[36] David F. Manlove. 2013. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company.

[37] Richard M. McCutchen. 2008. The Least-Unpopularity-Factor and Least-Unpopularity-Margin Criteria for Matching Problems with One-Sided Preferences. In *Proceedings of the 8th Latin American Conference on Theoretical Informatics (LATIN) (Lecture Notes in Computer Science (LNCS), Vol. 4957)*. 593–604.

[38] Eric McDermid and Robert W. Irving. 2011. Popular Matchings: Structure and Algorithms. *Journal of Combinatorial Optimization* 22, 3 (2011), 339–358.

[39] David C. McGarvey. 1953. A Theorem on the Construction of Voting Paradoxes. *Econometrica* 21, 4 (1953), 608–610.

[40] Kurt Mehlhorn and Dimitrios Michail. 2006. Network Problems with Non-Polynomial Weights and Applications. (2006). Unpublished manuscript. Available at https://pure.mpg.de/rest/items/item_2344037/component/file_2344036/content.

[41] Thomas Schwartz. 1986. *The Logic of Collective Choice*. Columbia University Press.

[42] Lloyd S. Shapley and H. Scarf. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1, 1 (1974), 23–37.

[43] John H. Smith. 1973. Aggregation of Preferences with Variable Electorate. *Econometrica* 41, 6 (1973), 1027–1041.

[44] Tayfun Sönmez and M. Utku Ünver. 2011. Matching, Allocation, and Exchange of Discrete Resources. In *Handbook of Social Economics*, Jess Benhabib, Matthew O. Jackson, and Alberto Bisin (Eds.). Vol. 1. Elsevier, Chapter 17, 781–852.

[45] Benjamin Ward. 1961. Majority rule and allocation. *Journal of Conflict Resolution* 5, 4 (1961), 379–389.