

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
Technical University of Munich
Munich, Germany
chris.dong@tum.de

Patrick Lederer
UNSW Sydney
Sydney, Australia
p.lederer@unsw.edu.au

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

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., 7, 17, 34, 41]. 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.

House allocation can be viewed as a special case of *social choice*, where agents have preferences over an abstract set of alternatives and a social choice function (SCF) returns the “best” alternative based on these preferences. 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. This perspective opens the door to transferring concepts from social choice theory to the domain of house allocation. For example, a natural way to compare two assignments μ and λ is to check whether a majority of agents prefer the house they receive under μ to the one they receive under λ .

Gärdenfors [26] pointed out that the resulting majority relation may be cyclic, demonstrating that a well-known issue in social choice—the *Condorcet paradox*—can also occur in the domain of assignment. Gärdenfors also introduced the notion of a *popular* assignment as an assignment, for which there is no other assignment that is preferred by a majority of the agents.¹ Popular assignments thus correspond to the well-studied notion of (weak) Condorcet winners in social choice theory and may fail to exist, just like in social choice. Popularity has been studied extensively and became an established concept in assignment [see, e.g., 2, 3, 9, 14, 18, 19, 36]. As majority cycles are even more frequent in assignment than they are in social choice, a crucial question is which assignment should be selected in the absence of popular assignments.² This has led to the definition of concepts such as *least unpopular* or *mixed popular* assignments [31, 35]. Interestingly, these assignment rules are also related to corresponding concepts studied in social choice theory: the *maximin* voting rule and *maximal lotteries*.

In this paper, we consider classic majoritarian SCFs in the context of assignment. An SCF is majoritarian if, like popularity, it is solely based on the majority relation, often viewed as a *majority graph*. Typical examples are Copeland’s rule, the top cycle, the uncovered set, and the bipartisan set. The definitions of most 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 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. When viewing house allocation as a subdomain of social choice, the number of alternatives is exponential, and the individual 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

We gain several insights that clarify the unique structure of the assignment domain, resulting in surprising outcomes that have no parallels in the broader social choice domain. We first characterize

¹Gärdenfors [26] referred to popular assignments as “majority assignments”.

²This can, for example, be seen by comparing the proportions of profiles that admit Condorcet winners and popular assignments, respectively [2, 27, Tables 1 and 4.2].

all sets of preference profiles that admit the same majority graph. This characterization is based on an efficient algorithm that reconstructs all profiles that induce a given majority graph. 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. Moreover, all Pareto-optimal assignments are contained in the top cycle, which means that elements of the top cycle can be found via serial dictatorships. Furthermore, 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 show 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. This characterization has no analog in the more general social choice domain, where the top cycle can consist of any number of alternatives, even when there are at most three agents.

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, 22].

Related Work

The study of matching under preferences was initiated by Gale and Shapley [25]. 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. Gärdenfors [26] showed that stable matchings are no longer guaranteed to exist when agents may be indifferent between other agents. He proposed to consider majority comparisons between matchings and introduced the notion of a popular matching, i.e., a matching such that there is another matching that a majority of the agents prefer. Popular matchings correspond to weak Condorcet winners in social choice theory. When individual preferences are strict, popularity is weaker than stability. However, popular matchings may not exist for weak preferences. This issue persists in the important special case where agents of one type are completely indifferent between all agents of the other type (these agents can then be seen as objects) and all other agents have strict preferences. This variant, which goes back to Gale [24], is known as assignment, house allocation, or two-sided matching with one-sided preferences. 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 [35] 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. [31] introduced mixed popular assignments, whose existence is guaranteed by the minimax theorem, and provided an efficient algorithm for computing them. Kavitha and Vaish [32] 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 hardness also holds in assignment. Kavitha and Vaish also give a fully polynomial-time randomized approximation scheme 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 [30]. Proposition 1 entails that semi-popular *assignments* can easily be found via serial dictatorships.

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*. This is a basic fairness notion, 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 single-valued assignment rule can be symmetric. We, therefore, consider set-valued rules. To nonetheless distinguish more 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 has been shown that an assignment is Pareto-optimal if and only if it is returned by a serial dictatorship for the given profile [1]. Pareto-optimality is much more restrictive in assignment than in the general social choice domain: 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., $\mu \succeq \lambda$ 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 [see, e.g., 8, 33]. Each profile P induces a directed graph $G_P = (M, \{(\mu, \lambda) \in M^2 : \mu \succeq \lambda\})$, 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. An arrow from μ to λ indicates that μ strictly majority dominates λ , and the absence of an arrow indicates a majority tie.



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 [37], 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.. 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 $H_{j+r} := H_{j+r-k}$ if $j+r > k$).

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

1: a, b, c, d	1: a, b, c, d
$\hat{P} =$ 2: a, b, c, d	$\bar{P} =$ 2: a, b, d, c
3: a, b, d, c	3: a, b, d, c
4: a, b, d, c	4: a, b, d, c

The profile $\bar{P}$ is *not* rotation equivalent to $\hat{P}$, as $c \succ_2 d$ and $d \succ'_2 c$ 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 all agent preferences except for the direction of some Pareto-dominations. The full proof of the following result is deferred to 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, q is added whenever not all agents prefer p to q or vice versa. By the insights of the previous paragraph, this means that we can determine for each agent which of the houses she prefers 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. $\square$

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. On the other hand, 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., 8, 33]. 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 $\mu \succ \lambda$ 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 [35], Kavitha et al. [31], and Brandt and Bullinger [9].

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 chooses the same set of assignments on both profiles.

Least unpopularity and mixed popularity: In the proof of Theorem 1, we show that we can infer the precise margin of each majority comparison from the majority graph. Since least popularity and mixed popularity only depend on these margins, they are majoritarian. $\square$

By contrast, we show next via an example that the rule that returns all *rank-maximal assignments* [29] is not majoritarian. To introduce rank-maximality, we define the rank of a house p in a preference relation $\succ$ by $r(\succ, p) = 1 + |\{q \in H \mid q \succ p\}|$, i.e., an agent’s favorite house has rank 1, his second favorite house has rank 2, etc. 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 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' , which are rotation equivalent and thus induce the same majority graph.

$P =$	1: a, b, c, d, e, f	$P' =$	1: d, e, f, a, b, c
	2: a, c, b, d, f, e		2: d, f, e, a, c, b
	3: b, a, c, e, d, f		3: e, d, f, b, a, c
	4: a, b, c, d, e, f		4: d, e, f, a, b, c
	5: a, b, c, d, e, f		5: d, e, f, a, b, c
	6: a, b, c, d, e, f		6: d, e, f, a, b, c

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 a surprising relationship between two majoritarian assignment rules that does not hold in the social choice domain. 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. Let a profile P be given. More formally, an assignment is *semi-popular* if $|\{\lambda \in M \mid \mu \succeq \lambda\}| \geq \frac{|M|}{2}$ [32]. By $SP(P)$, we denote the set of all semi-popular assignments.

Proposition 1. $PO \subseteq SP$.

PROOF. For our proof, we first introduce permutations on assignments. Given a permutation π on N , define $\pi': (N \rightarrow H) \rightarrow (N \rightarrow H)$ such that for any assignment $\mu: N \rightarrow H$ 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 preference 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 $|\{\lambda \in M \mid \mu \succeq \lambda\}| \geq \frac{|M|}{2}$, so μ is semi-popular. $\square$

The proof of Proposition 1, in fact, shows a stronger statement: each Pareto-optimal assignment has at least as many outgoing (strict) majority edges as ingoing ones.

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., 4, 15, 28, 39, 40]. The underlying idea is very natural: popular assignments do not always exist because the majority relation $\succeq$ fails to be transitive (see Example 1). Instead, one can consider $\succeq^*$, the *transitive closure* of $\succeq$, and simply return the maximal elements according to this relation. Formally, $TC(P) = \{\mu \in M : \forall \lambda \in M : \mu \succeq^* \lambda\}$.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. We defer the proof of the following result to Appendix B.1.

Proposition 2. $PO \subseteq TC$.

Using Proposition 2 as a stepping stone, we obtain a much stronger structural result about majority graphs in assignment: the top cycle can only contain one, two, all but one, all but two, all but three, or all assignments. The full proof of this result is given in Appendix B.

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 OF THEOREM 2. 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.

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 $n \geq 6$. For case (v), we introduce the concept of the bottom cycle, which contains all assignments that can be reached from every other assignment via a path in the majority graph. Using our base case for $n = 5$, we prove case (v) for $n \geq 6$ by carefully constructing

³One 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 [39].

The top cycle is unrelated to the *Top Trading Cycle (TTC)*, an assignment algorithm for settings with initial endowments.

⁴This 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.

an assignment μ from the top cycle and an assignment λ from the bottom cycle such that $\lambda \succ^* \mu$. This implies that the top cycle contains every assignment since every assignment has a path to every other assignment in the majority graph via λ and μ .

Lastly, we turn to the cases (iii) and (iv). For these cases, we can assume that the profile takes one of the following forms as otherwise case (i) or (ii) would apply.

$$\begin{array}{lll}
 P = \begin{array}{l} 1: a, \dots \\ 2: a, \dots \\ 3: a, \dots \\ 4: \dots \dots \\ 5: \dots \dots \end{array} & P' = \begin{array}{l} 1: a, c \\ 2: a, c \\ 3: b, d \\ 4: b, d \\ 5: \dots \dots \end{array} & P'' = \begin{array}{l} 1: a, b \\ 2: a, \dots \\ 3: b, \dots \\ 4: \dots \dots \\ 5: \dots \dots \end{array}
 \end{array}$$

We will exemplarily discuss our proof idea for the profile P'' . Contrary to the proof of case (v), we have no choice but to start from an arbitrary assignment λ^* not contained in the bottom cycle. For this assignment, we show that there is a trading cycle in which all agents involved get worse off. For example, the trading cycle may be of length two involving agents 4, 5 and houses c, d . Since all the agents in this trading cycle are worse off, we can let agent 1 swap her house with whoever currently holds house a . This yields an assignment λ with $\lambda^* \succ^* \lambda$ and $\lambda(1) = a$. We illustrate this step with the subsequent profile, where λ is the blue assignment and λ^* the red one.

$$\begin{array}{l}
 1: a, b, \dots \\
 2: a, \dots \dots \\
 P'' = 3: b, \dots \dots \\
 4: \dots c, d \\
 5: \dots d, c
 \end{array}$$

Take now any Pareto-optimal assignment μ^* with $\mu^*(1) = a$ and $\mu^*(3) = b$. We will define a sequence of assignments $\lambda^0, \dots, \lambda^k$ such that $\lambda^0 = \lambda$, $\lambda^k = \mu^*$, and $\lambda^0 \succ^* \lambda^1 \succ^* \dots \succ^* \lambda^k$. As a first step, we will construct an assignment λ^1 such that $\lambda^1(1) = a$ and $\lambda^1(3) = b$. Let N' be the set of 5 agents containing agents 1, 2, 3, the agent x with $\lambda^0(x) = b$, and an arbitrary fifth agent y . Moreover, let H' be the houses assigned to these agents in λ^0 . We define λ^1 as the assignment obtained from λ^0 by letting agents 3 and x swap their houses. When considering the restriction of P'' to N' and H' , we see that case (iii), (iv), or (v) apply to this 5-agent instance. From this, we infer that there is a path in the majority graph from λ^0 to λ^1 .

Finally, starting from λ^1 , we can repeatedly identify pairs of agents x and y such that agent y obtains in the current assignment λ^i the house $\mu^*(x)$. Then, we obtain the next assignment λ^{i+1} by letting agents x and y swap their assigned houses. By restricting the profile P'' to the agents $\{1, 2, 3, x, y\}$ and their currently assigned houses, we can again use the insights of the 5-agent case to deduce that there is a path from λ^i to λ^{i+1} . Finally, by repeatedly applying this argument, we get that there is a path from λ^* to μ^* in the majority graph, which shows that λ^* is in the top cycle. $\square$

Theorem 2 shows that deterministic assignments are highly unstable with respect to majority deviations. One can almost always, given a starting assignment and a target assignment, convince the agents to transition from one to the other by presenting intermediate assignments that are preferred by a weak majority of agents. To

illustrate this point, consider the following poor assignment that reaches every other assignment via some majority path.

Example 4. In P , the assignment μ marked in red is obviously not a desirable assignment. It fails to be Pareto-optimal, and many agents even receive their last-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 Appendix C. The rough idea is to reassign to some agents slightly worse houses while improving other agents' assignments significantly. This creates majority ties with better and better assignments. 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. These are now significantly happier.

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

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, it is possible that the top cycle contains $n! - 2$ 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.

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

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.

$$\begin{array}{ll}
 \begin{array}{l} 1: a, b, \textcolor{red}{c}, d \\ 2: a, b, \textcolor{red}{d}, c \\ 3: c, d, a, \textcolor{red}{b} \\ 3: c, d, b, \textcolor{red}{a} \end{array} & \begin{array}{l} 1: a, b, \textcolor{red}{c}, d \\ 2: a, b, \textcolor{red}{d}, c \\ 3: d, c, a, \textcolor{red}{b} \\ 3: d, c, b, \textcolor{red}{a} \end{array}
 \end{array}$$

Remark 4. Theorem 2 stands in stark contrast to classic social choice, where the top cycle has virtually no cardinality restrictions. In social choice, 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 [8, 12]. 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

We now proceed 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. Uncovered sets have been extensively studied in social choice theory [see, e.g., 5, 6, 11, 20].

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., $UC(P) = \{\mu \in M : \text{no } \lambda \in M \text{ covers } \mu\}$. 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 $UC \subseteq PO$ [23]. 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.

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

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 uncovered assignments set can be seen as particularly attractive among 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 agents preferences uniformly at random, depicted in Figure 2. It turns out that the Bordes-UC is almost 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




Figure 1: Size distributions of UCs for $n = 5$. The high peak is at size 2 for all of them. In total, there are 9078630 profiles (up to symmetries) and there are $5! = 120$ different assignments.




Figure 2: Size distributions of UCs for $n = 7$ in 1000 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.

such that $\lambda \succ \eta \succ \mu$. However, for small numbers of agents, there are many profiles admitting popular assignments. If μ is such a popular assignment, then 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 [8, 12]. In assignment, the uncovered set can be exponentially large. Unless a structural result such as




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, as percentage. In total, there are 9078630 profiles because we considered them up to symmetry. Note that the plots for McKelvey and Bordes almost perfectly align.

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, we could obtain an element of the uncovered set by computing a rank-maximal assignment, 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, yet contains all rank-maximal and all popular assignments.

Remark 7. Generous assignments [34, 38] 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

1:	<u>c</u>	f	a	e	b	g	d
2:	<u>b</u>	c	g	e	a	d	f
3:	<u>g</u>	f	a	<u>d</u>	e	c	b
$P =$ 4:	<u>g</u>	b	e	c	a	f	<u>d</u>
5:	<u>e</u>	d	a	b	c	f	g
6:	<u>a</u>	b	d	g	f	e	c
7:	<u>f</u>	b	d	e	c	a	g

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, somewhat surprisingly, revealed that some well-known assignment rules are in fact 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 (symmetric) 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., 8, 10, 21]. 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

⁵Kavitha and Vaish [32] have studied semi-popularity and Copeland winners in the more general setting of roommate markets.

investigate whether—in contrast to general 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 structural and computational properties.

ACKNOWLEDGMENTS

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/11-2. We greatly appreciate the contributions of Martin Bullinger and Matthias Greger during our initial research meetings on this topic.

REFERENCES

- [1] Atilla 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] Georges Bordes. 1976. Consistency, Rationality and Collective Choice. *Review of Economic Studies* 43, 3 (1976), 451–457.
- [5] Georges Bordes. 1983. On the Possibility of Reasonable Consistent Majoritarian Choice: Some Positive Results. *Journal of Economic Theory* 31, 1 (1983), 122–132.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] Felix Brandt and Martin Bullinger. 2022. Finding and Recognizing Popular Coalition Structures. *Journal of Artificial Intelligence Research* 74 (2022), 569–626.
- [10] Felix Brandt and Felix Fischer. 2008. Computing the Minimal Covering Set. *Mathematical Social Sciences* 56, 2 (2008), 254–268.
- [11] Felix Brandt and Felix Fischer. 2008. On the Hardness and Existence of Quasi-Strict Equilibria. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT) (Lecture Notes in Computer Science (LNCS), Vol. 4997)*. Springer-Verlag, 291–302.
- [12] 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.
- [13] Felix Brandt, Johannes Hofbauer, and Martin Suderland. 2016. Majority Graphs of Assignment Problems and Properties of Popular Random Assignments. In *Proceedings of the 6th International Workshop on Computational Social Choice (COMSOC)*.
- [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] 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.
- [18] Ágnes Cseh. 2017. Popular Matchings. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 6.
- [19] 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.
- [20] John Duggan. 2013. Uncovered Sets. *Social Choice and Welfare* 41, 3 (2013), 489–535.
- [21] Bhaskar Dutta and Jean-François Laslier. 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare* 16, 4 (1999), 513–532.
- [22] Mark Fey. 2008. Choosing From a Large Tournament. *Social Choice and Welfare* 31, 2 (2008), 301–309.
- [23] Peter C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489.
- [24] David Gale. 1960. *The Theory of Linear Economic Models*. McGraw-Hill.
- [25] David Gale and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.
- [26] Peter Gärdenfors. 1975. Match Making: Assignments based on bilateral preferences. *Behavioral Science* 20, 3 (1975), 166–173.
- [27] William V. Gehrlein and Peter C. Fishburn. 1979. Proportions of Profiles with a Majority Candidate. *Computers & Mathematics with Applications* 5, 2 (1979), 117–124.
- [28] I. J. Good. 1971. A Note on Condorcet Sets. *Public Choice* 10, 1 (1971), 97–101.
- [29] 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.
- [30] 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.
- [31] Telikepalli Kavitha, Julián Mestre, and M. Nasre. 2011. Popular mixed matchings. *Theoretical Computer Science* 412, 24 (2011), 2679–2690.
- [32] 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.
- [33] Jean-François Laslier. 1997. *Tournament Solutions and Majority Voting*. Springer-Verlag.
- [34] David F. Manlove. 2013. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company.
- [35] 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.
- [36] Eric McDermid and Robert W. Irving. 2011. Popular Matchings: Structure and Algorithms. *Journal of Combinatorial Optimization* 22, 3 (2011), 339–358.
- [37] David C. McGarvey. 1953. A Theorem on the Construction of Voting Paradoxes. *Econometrica* 21, 4 (1953), 608–610.
- [38] 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.
- [39] Thomas Schwartz. 1986. *The Logic of Collective Choice*. Columbia University Press.
- [40] John H. Smith. 1973. Aggregation of Preferences with Variable Electorate. *Econometrica* 41, 6 (1973), 1027–1041.
- [41] 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.

A PROOF OF THEOREM 1

In this appendix, we present the proof of Theorem 1, which we break down into the following lemmas.

Lemma 1. *Let $x, y \in N$ and $p, q \in H$ be pairwise distinct. Consider any μ with $\mu(x) = p$, $\mu(y) = q$. Consider the matching λ which only differs from μ in $\lambda(x) = q$ and $\lambda(y) = p$. Then, the following holds:*

- $\mu \succ \lambda$ iff $p \succ_x q$ and $q \succ_y p$,
- $\mu \prec \lambda$ iff $q \succ_x p$ and $p \succ_y q$, and
- $\mu \sim \lambda$ iff x, y have the same preferences over p, q (i.e., either $p \succ_x q$ and $p \succ_y q$ or $q \succ_x p$ and $q \succ_y p$).

PROOF. Since $\mu(z) = \lambda(z)$ for all agents $z \in N \setminus \{x, y\}$, the majority comparison only depends on the preferences of x, y between p, q . Going through all three cases, we see that the statement clearly holds. As aid, we visualize the assignments μ and λ in blue and red, respectively, for all three cases.

- x prefers p to q but y does not (hence $\mu \succ \lambda$):

$$P = \begin{array}{l} x: \quad p, \quad q \\ y: \quad q, \quad p \end{array}$$

- x prefers q to p , but y does not (hence $\mu \prec \lambda$):

$$P' = \begin{array}{l} x: \quad q, \quad p \\ y: \quad p, \quad q \end{array}$$

- x, y have the same preferences over p, q (hence $\mu \sim \lambda$):

$$P'' = \begin{array}{l} x: \quad q, \quad p \\ y: \quad q, \quad p \end{array} \quad P''' = \begin{array}{l} x: \quad p, \quad q \\ y: \quad p, \quad q \end{array}$$

□

Fix any $p, q \in H$. We apply Lemma 1 for the agent pair 1, 2, then 2, 3, and so on. Either, all agents have the same preferences over the pair p, q , or for some pair $x, x+1$ we can determine the preferences of these two agents over p, q and hence obtain these pairwise preferences of all agents.

We build a graph with the houses as nodes and edges between two assignments if we were able to fully determine the pairwise agent preferences between this pair. Name the connected components of this graph $H_1, \dots, H_k$. Next, we focus on each connected component H_i . We are able to fully determine the preferences within each component by virtue of the following lemma.

Lemma 2. *Let $p, q, r \in H$ be distinct such that for all agents, we know their preferences between p and q , as well as between q and r . Then, we also know all agents' preferences between p and r .*

PROOF. If not all agents rank p over r or not all agents rank r over p , we can apply Lemma 1 to obtain all agent preferences between p and r . Hence, let all agents rank p over r , or let all agents rank r over p . It now suffices to determine the preferences of a single agent. If there is an agent x such that $p \succ_x q \succ_x r$ or $r \succ_x q \succ_x p$, then by transitivity we know $p \succ_x r$ or $r \succ_x p$, respectively. Otherwise, all agents x rank q either first or last among $\{q, p, r\}$ (i.e., either $q \succ_x p, r$ or $p, r \succ_x q$). Because the case $n \leq 2$ is trivial, consider $n \geq 3$. By the pigeonhole principle, there are at least two agents x and y of the same type. We only consider the case where $q \succ_x p, r$ and $q \succ_y p, r$. The other case is symmetric and follows from the same arguments. Now, take any third agent $z \in N$ and consider the two assignments μ and λ which we first visualize.

$$P = \begin{array}{l} x: \quad q, \quad p, \quad r \\ y: \quad q, \quad p, \quad r \\ z: \quad \dots \quad p, \quad r \end{array}$$

Formally, we take any $\mu \in M$ with $\mu(x) = q, \mu(y) = p$, and $\mu(z) = r$. We obtain λ by altering μ in $\lambda(x) = r, \lambda(y) = q$, and $\lambda(z) = p$. Since $\mu(w) = \lambda(w)$ for all $w \in N \setminus \{x, y, z\}$, the majority comparison only depends on the agents x, y, z . Clearly, x prefers μ to λ , and y prefers λ to μ . Hence, there are two possibilities. If $\mu \succ \lambda$, then $r \succ_z p$. Otherwise $\mu \prec \lambda$, due to $r \prec_z p$. We have thus determined the preferences of agent z regarding houses p and r . This concludes the proof. □

Based on the last two lemmas, we can now identify sets of houses $H_1, \dots, H_k$ for which we can fully specify the agents' preferences. To this end, consider the following graph G , whose vertices are the houses H , and there is an edge between two houses p and q if applying Lemma 1 for all agent pairs allows us to infer the agents' preferences over p and q . Then, the sets $H_1, \dots, H_k$ correspond to the connected components of this graph. In particular, for any two houses p, q that are connected in this graph, there is a sequence of houses $p_1 = p, \dots, p_\ell = q$ such that p_i is a neighbor of p_{i+1} in G . By repeatedly applying Lemma 2 along this path, we can infer the preferences of all agents between p and q .

Moreover, we claim that the houses in each set H_i must be ranked contiguously by each agent, i.e., for all H_i , agents $x \in N$, and houses $p, q \in H_i$, there is no house $r \notin H_i$ such that $p \succ_x r \succ_x q$. Assume for contradiction that this is not true, which means that there is a connected component H_i of G , an agent x , and houses $p, q \in H_i$, $r \notin H_i$ such that $p \succ_x r \succ_x q$. Further, we partition H_i into the sets $H_i^+ = \{p' \in H_i : p' \succ_x r\}$ and $H_i^- = \{q' \in H_i : r \succ_x q'\}$. Since $r \notin H_i$, we know that there is no edge from r to any house $h \in H_i$ in G . This means all agents agree on the preference between r and the houses in H_i , i.e., it holds for all $y \in N$ that $p' \succ_y r$ for all $p' \in H_i^+$ and $r \succ_y q'$ for all $q' \in H_i^-$. However, by the transitivity of the agents' preferences, this means that $p' \succ_y q'$ for all $y \in N$, $p' \in H_i^+$, and $q' \in H_i^-$. In turn, this implies that there is no edge between houses in H_i^+ and H_i^- in G , which contradicts that H_i is a connected component in this graph. Note that contiguity of each component H_i is only required for correctness of the algorithm, not for the algorithm itself.

We note that it follows from the last paragraph also that, for all components H_i, H_j , it either holds that $p \succ_x q$ for all $x \in N$, $q \in H_i$, $p \in H_j$, or $q \succ_x p$ for all $x \in N$, $q \in H_i$, $p \in H_j$. The reason for this is that all agents rank the houses in H_i and H_j contiguously. Moreover, if there were two agents x and y such that $p \succ_x q$ and $q \succ_y p$ for all $p \in H_i$, $q \in H_j$, then H_i and H_j would be placed in the same connected component. From this, we infer that the sets $H_1, \dots, H_k$ can be ordered to form a valid decomposition of P .

By our analysis so far, we get that every profile that induces G_P can only differ in the order of the components $H_1, \dots, H_k$ in the agents' preferences. To complete the proof, we need to show that the only cyclic shifts of the components lead to the same majority graph. This last step can be inferred analogously to Lemma 4 of Brandt et al. [13]. The idea is that, for each three distinct houses p, q, r , a linear preference relation either agrees with two out of three of the following comparisons $p \succ_{cyc} q, q \succ_{cyc} r, r \succ_{cyc} p$,

or it agrees with two out of three of the following comparisons $p \succ'_{cyc} r, r \succ'_{cyc} q, q \succ'_{cyc} p$. We can use this as follows:

Lemma 3. *Let p, q, r be houses such that agents 1, 2, 3 satisfy $\succ_1 \mid_{p,q,r} = \succ_2 \mid_{p,q,r} = \succ_3 \mid_{p,q,r}$. Then, from the majority graph, we can determine the “direction” in which these three are ordered.*

PROOF. Consider any μ with $\mu(1) = p, \mu(2) = q, \mu(3) = r$, and λ which differs from μ only in $\lambda(1) = q, \lambda(2) = r, \lambda(3) = p$. Then, it holds that $\mu \succ \lambda$ or $\lambda \succ \mu$, depending on whether the preferences are of type $p \succ'_{cyc} r, r \succ'_{cyc} q, q \succ'_{cyc} p$ or $p \succ_{cyc} q, q \succ_{cyc} r, r \succ_{cyc} p$, respectively. $\square$

For each three distinct components, we can take $p \in H_i, q \in H_j, r \in H_\ell$ and apply Lemma 3. Since each component is contiguous in all preference relations, this gives us the “cycle type” of how the three components are ordered in the entire profile. Now, one can prove that any cyclic permutation of the components does not change the “cycle type” of any three distinct components. However, any non-cyclic permutation of the components yields three distinct components H_i, H_j, H_r such that their cycle type is changed. By choosing $p \in H_i, q \in H_j, r \in H_r$, we can use the assignment in Lemma 3 to obtain μ and λ for which the majority comparison does not coincide with the respective one in G_{P^*} .

B PROOF OF THEOREM 2

In this section, we present the proof of Theorem 2 in two steps. First, we introduce Pareto-pessimality and the bottom cycle, which are dual concepts to Pareto-optimality and the top cycle. Then, we prove that all Pareto-optimal assignments are contained in the top cycle and characterize the cases in which the top cycle is of size at most two. We obtain analogous statements for the bottom cycle. Then, in step two, we show that the top cycle and bottom cycle coincide whenever they are not of size one or two. This requires a proof by induction. We prove the base case for five agents by computer, whereas the induction step is shown by hand.

For the proof of Theorem 2, we have to consider concepts for assignments that are particularly “bad”. Given a profile P , by P^{-1} , we denote the profile where all agent preferences $\succ_x$ are inverted to $\succ_x^{-1}$, i.e., $p \succ_x q$ if and only if $q \succeq_x^{-1} p$. The majority relation induced by P^{-1} is precisely the inverse of the majority relation $\succeq$ of P , and we hence denote it by $\succeq^{-1}$.

As a dual concept to a serial dictatorship, a serial antidictatorship works exactly the other way around. The agents pick their least preferred houses that is available yet in order of σ .

Analogously to the TC , the bottom cycle consists of all assignments which are minimal elements in the transitive closure of the weak dominance relation. Formally, $BC(P) = \{\mu \in M_{N,H} : \forall v : v \succeq^* \mu\}$. An assignment is *Pareto-pessimal* if it does not Pareto-dominate any other assignment. By $PP(P) = PO(P^{-1})$, we denote the set of all Pareto-pessimal assignments in P . It holds that an assignment is Pareto-pessimal if and only if it can be obtained as a serial antidictatorship.

Lemma 4. *For all profiles P , we have $TC(P) = BC(P^{-1})$. (Note that this directly implies $BC(P) = BC((P^{-1})^{-1}) = TC(P^{-1})$.)*

PROOF. Note that the majority relation over the assignments induced by P is precisely inverse to the majority relation induced

by P^{-1} , and hence so is the transitive closure over it. Thus, $\mu \succeq^* \lambda$ for all assignments λ if and only if $\lambda (\succeq^{-1})^* \mu$ for all assignments λ . This proves the claim. $\square$

For example, Lemma 4 implies that each Pareto-pessimal assignment belongs to the bottom cycle, as Pareto-pessimal assignments of a profile P are Pareto-optimal assignments of P^{-1} .

B.1 Step 1: Understanding Small TC sizes

We first characterize all profiles in which TC chooses at most two assignments. For this, we prove a strengthening of Proposition 2.

Proposition 3. *All Pareto-optimal assignments are in the top cycle.*

$$PO \subseteq TC$$

Conversely, all Pareto-pessimal assignments are in the bottom cycle.

$$PP \subseteq BC$$

PROOF. It suffices to show $PO \subseteq TC$, as the proof for $PP \subseteq BC$ then follows from Lemma 4. For this, we will prove that, for every profile P , each Pareto-optimal assignment can reach each other Pareto-optimal assignment via a path in the majority graph. This shows that $PO(P) \subseteq TC(P)$ because each assignment that is not Pareto-optimal is Pareto-dominated (and thus also majority dominated) by an assignment in $PO(P)$. To prove this claim, we will rely on the characterization of Pareto-optimal assignments by Abdulkadiroğlu and Sönmez [1], which states that an assignment μ is in $PO(P)$ if and only if there is a priority order σ over the agents such that μ is the outcome of the corresponding serial dictatorship.

Before showing that all Pareto-optimal assignments are connected by paths in the majority graph, we prove an auxiliary statement: when modifying a picking sequence σ by improving the position of the last agent, the outcome of the serial dictatorship for the original sequence weakly majority dominates the outcome of the serial dictatorship for the modified sequence. To make this more formal, fix an order $\sigma = (x_1, \dots, x_n)$ over the agents and let $\sigma' = (x_1, \dots, x_{k-1}, x_n, x_k, \dots, x_{n-1})$ for some $k \in \{1, \dots, n-1\}$. Moreover, let μ denote the assignment picked by the serial dictatorship induced by σ and λ the one picked by the serial dictatorship induced by σ' . We will show that $\mu \succeq \lambda$.

To prove this, let X_i and X'_i denote the houses that are available when agent x_i gets to pick her house under σ and σ' , respectively. We claim that $X_i \supseteq X'_i$ for all $i \in \{1, \dots, n-1\}$. First, for the agents $x_i \in \{x_1, \dots, x_{k-1}\}$, it holds even that $X_i = X'_i$ because σ and σ' agree on the first $k-1$ agents. Next, we have that $X_k = X'_k \cup \{\lambda(x_n)\}$ because agent x_k gets to pick before x_n in σ . Now, inductively assume that $X_i \supseteq X'_i$ for some $i \in \{k, \dots, n-2\}$. We will show that $X_{i+1} \supseteq X'_{i+1}$. For this, we note that $|X_i| = |X'_i| + 1$ because the agents $x_1, \dots, x_{i-1}$ pick before x_i under σ , whereas x_n additionally gets to choose before x_i under σ' . By our induction assumption, we thus conclude that there is a single house p such that $X_i = X'_i \cup \{p\}$. Since $\lambda(x_i)$ is agent x_i 's favorite house in X'_i , this means that agent x_i either picks p or $\lambda(x_i)$ from X_i . Consequently, $X_{i+1} = X'_i$ if agent x_i picks p , or $X_{i+1} = X'_{i+1} \cup \{p\}$ if she picks $\lambda(x_i)$. In both cases, it holds that $X_{i+1} \supseteq X'_{i+1}$, thus proving the induction step.

By the definition of serial dictatorships, the fact that $X_i \supseteq X'_i$ for all $x_i \in \{x_1, \dots, x_{n-1}\}$ implies that $\mu \succeq_{x_i} \lambda$ for all these agents. If one of these agents strictly prefers μ to λ , we have that $\mu \succeq \lambda$

as only agent x_n may prefer λ to μ . On the other hand, if no $x_i \in \{x_1, \dots, x_{n-1}\}$ strictly prefers μ to λ , then $\mu(x_i) = \lambda(x_i)$ for all these agents, which implies that $\mu = \lambda$ and thus again $\mu \succeq \lambda$.

Based on our auxiliary claim, we will now complete the proof of this lemma. To this end, we fix two distinct assignments $\mu, \lambda \in PO(P)$ and consider two orders $\sigma = (x_1, \dots, x_n)$ and $\sigma' = (x'_1, \dots, x'_n)$ such that the corresponding serial dictatorship choose μ and λ for P . We will iteratively transform σ into σ' . To this end, let i denote the smallest index such that $x_i \neq x'_i$ and let $j > i$ denote the index such that $x_j = x'_i$. Now, consider the priority order σ_1 derived from σ by placing the currently last-ranked agent directly before x_j , and let η_1 be the assignment chosen by the corresponding serial dictatorship. By our auxiliary claim, we have $\mu \succeq \eta_1$. Further, we can repeat this step until we have a priority order σ_2 where $x'_i = x_j$ is in the last position, and our auxiliary claim shows that $\mu \succeq^* \eta_2$ for the assignment η_2 chosen by the corresponding assignment. Next, let σ_3 be the priority order derived from σ_2 by moving $x_{i'}$ directly before x_i . Again using our auxiliary claim, the assignment η_3 chosen by the corresponding serial dictatorship satisfies that $\eta_2 \succeq \eta_3$, so we have $\mu \succeq^* \eta_3$. Lastly, we observe that the priority orders σ_3 and σ' agree on the first i positions. Hence, by repeating our argument at most n times, we transform σ_1 into σ_2 while constructing a path from μ to λ in the majority graph. $\square$

In the remainder of this subsection, we use Proposition 3 to analyze cases (i) and (ii) of Theorem 2, as well as analogous statements regarding the bottom cycle.

We next use Proposition 3 to analyze profiles where all agents have distinct top choices (or bottom choices).

Lemma 5. $|TC(P)| = 1$ if and only if all agents have distinct top choices.

Analogously, $|BC(P)| = 1$ if and only if all agents have distinct bottom choices.

PROOF. Let P be a profile where all agents have distinct top choices. Let μ be the assignment that assigns to each agent her top choice. For all other assignments λ , there is at least one agent who does not obtain their top choice. These agents strictly prefer μ to λ , while all other agents weakly prefer μ to λ . Thus, we have $\mu \succ \lambda$ for all assignments $\lambda \neq \mu$, which implies $TC(P) = \{\mu\}$.

We prove the other implication for TC by contraposition. Let P be a profile where two agents $x \neq y$ have the same favorite house p . By Proposition 3, we know that all assignments obtained via serial dictatorships are contained in the top cycle. Since the serial dictatorship where x chooses first and y second yields a different assignment than the serial dictatorship where y chooses first and then x chooses, there are at least two assignments in the top cycle. This proves the statement for the top cycle.

For the bottom cycle, the statement follows from Lemma 4. This concludes the proof. $\square$

The first part of Lemma 5 follows from existing literature (e.g., 34, Proposition 7.24.). For the sake of completeness, we provided a proof nevertheless. We further note that Lemma 5 directly proves case (i) of Theorem 2 and will be useful for the proof of case (iv).

Next, we describe all cases in which the top cycle has cardinality two. We analogously obtain a similar statement for bottom cycle.

Lemma 6. $|TC(P)| = 2$ if and only if all but two agents have distinct top choices and the two agents who have the same top choice also share the same second best choice, which is not top-ranked by any agent.

Analogously, $|BC(P)| = 2$ if and only if all but two agents have distinct bottom choices and the two agents who have the same bottom choice also share the same second worst choice, which is not bottom-ranked by any agent.

PROOF. Let P be a profile where two distinct agents x^*, y^* share the same top choice p and second choice q , and all other agents have unique favorite houses in $H \setminus \{p, q\}$. Consider the two assignments μ_{x^*} , where all agents but x^* obtain their favorite house and x^* obtains q , and μ_{y^*} , where all agents but y^* obtain their favorite houses and y^* obtains q . Clearly, $\mu_{x^*} \sim \mu_{y^*}$. Further, we note that in every other assignment $\lambda \in M \setminus \{\mu_{x^*}, \mu_{y^*}\}$, at least one agent among x^* and y^* cannot obtain their favp. Without loss of generality, let that agent be x^* . Then, every agent weakly prefers μ_{x^*} to λ . Moreover, since $\lambda \notin \{\mu_{x^*}, \mu_{y^*}\}$, there is at least one agent $x \in N \setminus \{x^*, y^*\}$ with $\lambda(x) \neq \mu_{x^*}(x)$. Indeed, if λ would agree with μ_{x^*} on all agents in $N \setminus \{x^*, y^*\}$, then either x^* gets p and y^* gets q and $\lambda = \mu_{y^*}$, or x^* gets q and y^* gets p and $\lambda = \mu_{x^*}$. This agent x strictly prefers μ_{x^*} to λ , so $\mu_{x^*} \succ \lambda$. Further, we compare λ to μ_{y^*} . If $\lambda(y^*) \neq p$, we have that $\mu_{y^*} \succeq_x \lambda$ for all $x \in N$, so we immediately get that $\mu_{y^*} \succ \lambda$. On the other hand, if $\lambda(y^*) = p$, then $\mu_{y^*}(x^*) = p \succ_{x^*} \lambda(x^*)$. Further, since there is another agent $x \notin \{x^*, y^*\}$ with $\mu_{y^*} \succ_x \lambda$, there are at least two agents strictly preferring μ_{y^*} to λ , proving that $\mu_{y^*} \succ \lambda$. This concludes the proof that $TC(P)$ has cardinality two.

For the other direction of the TC statement, let P not be of the above form. Then, we distinguish between three cases.

- (1) The top choices of the agents are all pairwise different.
- (2) There are three agents who share the same top choice.
- (3) All agents but x^* and y^* have disjoint top choices, but agents x^* and y^* have different second choices.
- (4) There are two pairs of agents x, y and z, w with coinciding top choices.

In Case (i), Lemma 5 directly implies that the cardinality of TC is not equal to two. In Case (ii), we consider three serial dictatorships, where one of these agents chooses their house first, respectively. Since the three serial dictatorships yield different outcomes, we obtain by Proposition 3 that the top cycle contains at least three assignments. In Case (iii), we again consider three serial dictatorships. By pigeon hole principle, the second choice of one of the two agents x^*, y^* coincides with the top choice of some agent z . Without loss of generality, we assume that the favorite house of agent z is the second-ranked house of agent y^* . In the first serial dictatorship, y^* chooses before x^* before z^* . In the second, the roles of x^* and y^* are reversed. In the third, z chooses first, x^* second, and y^* third. The resulting three assignments do not coincide as agent y^* obtains her first-ranked (resp. second-ranked or third-ranked) house in the first (resp. second or third) serial dictatorship. Hence, in Case (iii) there are at least three assignments in the top cycle by Proposition 3. Finally, in Case (iv), we can consider all serial dictatorships which vary the picking order between agents x, y, z , and w . This results in at least four assignments that are serial dictatorships and hence in the top cycle. This concludes the proof of the top cycle statement.

Note that the bottom cycle statement follows from Lemma 4 and the proven statement about the top cycle. $\square$

Lemma 6 directly proves case (ii) of Theorem 2. It will further be helpful in proving case (iii).

We will make case analyses based on the structure of the profile. For often-recurring structures, we introduce shorthand notation. Let $\mathcal{P}^{p,p,p}$ denote the set of profiles where three distinct agents have p as their top-ranked house. We use concatenation to describe further ranks of agents, i.e., $\mathcal{P}^{p,q,p,q}$ denotes the set of profiles where two distinct agents have p as their top-ranked house and q as their second choice. Similarly, $\mathcal{P}_{p,p}$ denotes the set of profiles where two distinct agents have p as their last-ranked house, and $\mathcal{P}_{p,q,p,q}$ means that these two agents also have q as their second-to-last choice. Now that we understand when top cycle and bottom cycle have size one or two, we summarize and slightly reformulate Lemma 5 and Lemma 6 using our new notation.

Lemma 7. *Let P be given. The following equivalences hold:*

- $|TC(P)| = 1 \iff \nexists p : P \in \mathcal{P}^{p,p}$
- $|TC(P)| = 2 \iff \exists p \in H : \forall q \neq r \in H : P \in \mathcal{P}^{p,p} \setminus \{\mathcal{P}^{q,q,q} \cup \mathcal{P}^{q,q,r,r} \cup \mathcal{P}^{q,r,q,r}\}$
- $|TC(P)| > 2 \iff \exists p, q \in H : P \in \mathcal{P}^{p,p,p} \cup \mathcal{P}^{p,p,q,q} \cup \mathcal{P}^{p,q,p,q}$

Analogously,

- $|BC(P)| = 1 \iff \nexists p \in H : P \in \mathcal{P}_{p,p}$
- $|BC(P)| = 2 \iff \exists p \in H : \forall q, r \in H : P \in \mathcal{P}_{p,p} \setminus \{\mathcal{P}_{q,q,q} \cup \mathcal{P}_{q,q,r,r} \cup \mathcal{P}_{q,r,q,r}\}$
- $|BC(P)| > 2 \iff \exists p, q \in H : P \in \mathcal{P}_{p,p,p} \cup \mathcal{P}_{p,p,q,q} \cup \mathcal{P}_{p,q,p,q}$

Note that for $|TC(P)| > 2$, a finer case analysis of the different structures is: $P \in \mathcal{P}^{p,p,p} \cup \mathcal{P}^{p,q,p,q} \cup \mathcal{P}^{p,r,p,r,q,s}$ for some pairwise disjoint $p, q, r, s \in H$. (Analogous cases for the bottom cycle hold.)

PROOF. By Lemma 4, it suffices to prove the TC statements.

Clearly, all top choices being distinct is equivalent to no house being top-ranked twice, and hence the $|TC(P)| = 1$ statement follows directly from Lemma 5.

For the $|TC(P)| = 2$ statement, we prove both implications. Let any P be given where TC has cardinality two. Lemma 6 states that, then, (a) all agents but two have distinct top choices. Further, (b) these two agents have the same second rank house, which no agent ranks as their top. Clearly, this implies that no three agents have the same top choice and there are not two distinct houses top-ranked by two agents each. Further, if there are two agents who top rank some house q and one of these ranks some house r second, then both of these agents rank r second. Hence, no other agent may top rank the second-ranked house r of these two agents, and $P \notin \mathcal{P}^{q,r,q,r}$.

For the other implication, let P be given such that the top cycle has a cardinality not equal to two. Then, we can negate the structure described in Lemma 6. If (a) is false, then the first case is that there are only distinct top choices, and hence $P \notin \mathcal{P}^{p,p}$ for any house p . Otherwise, at least three agents top-ranking the same house p , or there exist at least two distinct houses that are top-ranked twice each. In this case, $P \in \mathcal{P}^{q,q,q} \cup \mathcal{P}^{q,q,r,r}$ for some houses $q \neq r$. If (a) is true but (b) is false, then the two agents sharing some top choice q either have the same second-ranked house r which is top-ranked

by some other agent—in which case we directly have $P \in \mathcal{P}^{q,r,q,r}$ —or they have distinct second-ranked houses r, s . There are $n - 2$ other agents, but only $n - 3$ houses not top or second-ranked by the two agents. The pigeonhole principle implies that one of these second-ranked houses, w.l.o.g. say r , is top-ranked by some agent. Hence, $P \in \mathcal{P}^{q,r,q,r}$.

For the $|TC(P)| > 2$ statement, we once more prove both directions. Note that if P is such a profile, then TC neither has cardinality one nor two. Hence, by our previously proven equivalences, there is some p' such that $P \in \mathcal{P}^{p',p'}$. Further, for this p' , there exist q', r' such that $P \in \mathcal{P}^{q',q',q'} \cup \mathcal{P}^{q',q',r',r'} \cup \mathcal{P}^{q',r',q',r'}$. By setting $p = q'$ and $q = r'$, we hence obtain $P \in \mathcal{P}^{p,p,p} \cup \mathcal{P}^{p,p,q,q} \cup \mathcal{P}^{p,q,p,q}$. For the other implication, let P be a profile such that $P \in \mathcal{P}^{p,p,p} \cup \mathcal{P}^{p,p,q,q} \cup \mathcal{P}^{p,q,p,q}$. Then, by definition of the three structures, $P \in \mathcal{P}^{p,p}$, proving that TC does not have cardinality one. Further, by setting $q' = p$ and $r' = q$, we obtain that $P \in \mathcal{P}^{q',q',q'} \cup \mathcal{P}^{q',q',r',r'} \cup \mathcal{P}^{q',r',q',r'}$, and hence TC does not have cardinality two either.

Finally, note that if $P \in \mathcal{P}^{q,q,q,r,r} \setminus \mathcal{P}^{q,s,q,s,r,r}$ for some $p \neq q \in H$ and all $s \neq t \in H$, then the first and second ranks of these four agents contain at least five houses. By the pigeon-hole principle, a fifth agent must top-rank a house which is among these four. If this house is q or r , then w.l.o.g. $P \in \mathcal{P}^{q,q,q}$. Otherwise, the house is the second rank s of some agent top ranking q or r , hence w.l.o.g. $P \in \mathcal{P}^{q,s,q,s}$. This concludes the proof. $\square$

Given Lemma 7, we call $TC(P)$ ($BC(P)$) *isolated* iff $|TC(P)| \leq 2$ ($|BC(P)| \leq 2$). Given a profile P and an assignment μ , we say that μ is not in an isolated top cycle, if $TC(P)$ is not isolated or if the top cycle is isolated and $\mu \notin TC(P)$.

B.2 Step 2: Understanding Large TC Sizes

By using Lemma 7, we can check for a given instance whether the top or bottom cycle are of size one or two. In contrast, at this moment of the proof, we know little about the exact size of the top cycle when it exceeds two. It will turn out that, if the top cycle contains more than two assignments, the transitive closure of the majority relation forms at most two equivalence classes: the top cycle and the bottom cycle. For now, we focus on the five-agent case, which will be very useful for later arguments.

Fact 1. *For any profile P with $n = 5$, the following statements hold.*

- (1) $|TC(P)| > 2 \iff M \setminus PP(P) \subseteq TC(P)$.
- (2) $|BC(P)| > 2 \iff M \setminus PO(P) \subseteq BC(P)$.
- (3) $|TC(P)| > 2$ and $|BC(P)| > 2 \iff TC(P) = M$
- (4) If $|TC(P)| \leq 2$ and $|BC(P)| \leq 2$, then for all $\mu, \lambda \in M \setminus (TC(P) \cup BC(P))$ we have $\mu \succeq^* \lambda$.

In words, if the top cycle (bottom cycle) contains more than two assignments, then it contains all assignments that are not serial anti-dictatorships (serial dictatorships). Further, if both contain at least three assignments, TC chooses all assignments. Vice versa, if both the top and bottom cycle contain two or fewer assignments, respectively, there is precisely one equivalence class sandwiched between the two.

PROOF. For $n = 5$, the number of possible instances is manageable: up to symmetry, there are about nine million preference profiles. Using Lemma 7, we are able to verify Fact 1 by brute forcing

through all possible profiles with the help of a household computer in two days. The code is in the supplementary material. $\square$

Fact 1 implies a very helpful lemma for general n . We first introduce corresponding terminology.

Restricted feasible sets. We now consider an arbitrary number of agents and require notation for restricted feasible sets. Recall that N, H induce a corresponding set of assignments M . Let $M_{N', H'}$ denote the set of all assignments on the agent set $N' \subseteq N$ and house set $H' \subseteq H$. Some assignments $\mu \in M$ also obtain a corresponding restriction $\mu_{N', H'}$, with domain N' and image H' . Note that we only use this notation if indeed $\mu(N') = H'$ is a bijection, and the restricted matching is hence well-defined on N' and H' . Similarly, for a profile P for agent set N and house set H , let $P_{N', H'}$ denote the profile restricted to agent set N' and house set H' . The restricted profile $P_{N', H'}$ induces a majority relation on $M_{N', H'}$. This relation is denoted via $\succeq_{N', H'}$. Note that we will only use this notation and compare two assignments w.r.t. $\succeq_{N', H'}$ if they both belong to $M_{N', H'}$.

Lemma 8 (Non-isolation). *Let $\mu, \lambda \in M$ such that for some N' of size 5, $\mu(N') = \lambda(N') =: H'$. If $\mu_{N', H'}$ is not part of an isolated top cycle w.r.t. $M_{N', H'}$, and $\lambda_{N', H'}$ is not part of an isolated bottom cycle w.r.t. $M_{N', H'}$, then $\lambda_{N', H'} \succeq_{N', H'}^* \mu_{N', H'}$. Moreover, if $\mu(x) = \lambda(x)$ for all $x \in N \setminus N'$, we have $\lambda \succeq^* \mu$.*

PROOF. If $TC(P_{N', H'})$ is not isolated and $BC(P_{N', H'})$ is neither, then $TC(P_{N', H'}) = M_{N', H'}$ by Claim (3) of Fact 1 and the statement follows trivially. If $TC(P_{N', H'})$ is not isolated but $BC(P_{N', H'})$ is, then $TC(P_{N', H'}) = M_{N', H'} \setminus BC(P_{N', H'})$. In more detail, our assumptions mean that $TC(P_{N', H'}) \neq BC(P_{N', H'})$, Claim (1) of Fact 1 shows that $M \setminus PP(P_{N', H'}) \subseteq TC(P_{N', H'})$, and Proposition 3 that $PP(P_{N', H'}) \subseteq BC(P_{N', H'})$, which implies that $TC(P_{N', H'}) = M_{N', H'} \setminus BC(P_{N', H'})$. Since λ is not in the isolated BC, $\lambda \in TC(P_{N', H'})$ and the lemma follows. Thirdly, assume that $TC(P_{N', H'})$ is isolated but $BC(P_{N', H'})$ is not. Then, we can infer from Claim (2) of Fact 1 and Proposition 3 that $TC(P_{N', H'}) = M_{N', H'} \setminus BC(P_{N', H'})$. Since $\mu_{N', H'}$ is not in the isolated top cycle, $\mu_{N', H'} \in BC(P_{N', H'})$ and the statement follows in this case. Finally, if $TC(P_{N', H'})$ is isolated and $BC(P_{N', H'})$ is too, then all assignments in $M_{N', H'} \setminus (TC(P_{N', H'}) \cup BC(P_{N', H'}))$ are connected via $\succeq_{N', H'}^*$ by Claim (4) of Fact 1. Since $\mu_{N', H'}$ is not in the isolated top cycle and $\lambda_{N', H'}$ is not in the isolated bottom cycle, we have $\mu_{N', H'} \in BC(P_{N', H'})$, or $\lambda_{N', H'} \in TC(P_{N', H'})$, or both belonging to the connected set of assignments. In all three cases, the statement follows directly. $\square$

Using Fact 1 and Lemma 8, we show in several steps that, whenever both top cycle and bottom cycle contain strictly more than two assignments, the two cycles coincide and contain all assignments.

Proposition 4. *Let P be any profile with $n \geq 5$. Then, $|TC(P)| > 2$ and $|BC(P)| > 2 \iff TC(P) = M = BC(P)$.*

PROOF. First, we quickly handle some easy cases. The direction from right to left is trivial as $n! > 2$. For $n = 5$, we obtain the direction from left to right directly from Fact 1. Now, let $n > 5$ and consider direction from left to right. Let any profile P be given such that $TC(P)$ and $BC(P)$ contain at least three assignments each.

In order to show $TC(P) = BC(P) = M$, it suffices to show that $\lambda^* \succeq^* \mu^*$ for some assignments $\mu^* \in TC(P)$ and $\lambda^* \in BC(P)$. To define these, we consider the structure of the profile.

Defining μ^* and λ^* . By Lemma 7, there are agents $x, y, z, w \in N$ and distinct houses $p, q, r, s \in H$ such that P has one of the following three structures.

Structure $p-p-p$: x, y, z are distinct and rank p first (set $w := z$), or

Structure $pq-p-q$: x, y, z are distinct, x, y rank p first, x ranks q second, and z ranks q first (set $w := z$), or

Structure $pr-pr-qs-qs$: x, y, z, w are distinct, x, y rank p first and r second, and z, w rank q first and s second.

Clearly, the profile $P_{N', H'}$ still contains the respective structure for all $N' \supseteq \{x, y, z, w\}$. To formally talk about the rank at which agents rank their houses, we define $h_x(i)$ to be the i th favorite house of agent x .

We aim to choose $\mu^* \in M$ with

- $\mu^* \in TC(P)$.
- $\mu^*(x) = p$.
- Under Structure $pr-pr-qs-qs$, we require $\mu^*(y) \neq h_y(2)$.

Under Structure $p-p-p$ and $pq-p-q$, this is achieved by letting $\mu^* \in PO$ with x picking first. Under Structure $pr-pr-qs-qs$, we first let μ be a serial dictatorship with priority order $\sigma = (y, x, z, \dots)$. Then, $\mu(x) = r, \mu(y) = p, \mu(z) = q$. We now permute to create μ^* . Let $\mu^*(x) = p, \mu^*(y) = q, \mu^*(z) = r$ and $\mu^*(v) = \mu(v)$ for all $v \in N \setminus \{x, y, z\}$. Restricted to $N' \supseteq \{x, y, z, w\}$ of size 5, μ and μ^* are not part of an isolated top or bottom cycle, respectively. By Lemma 8, $\mu \succeq^* \mu^*$, hence $\mu^* \in TC(P)$. We can hence choose μ^* as desired.

Analogously, since the bottom cycle contains more than two assignments, there are $x', y', z', w' \in N$ and distinct $p', q', r', s' \in H$ such that

Structure $p'-p'-p'$: x', y', z' are distinct and rank p' last (set $w' := z'$), or

Structure $p'q'-p'-q'$: x', y', z' are distinct, x', y' rank p' last, x' ranks q' second-to-last, and z' ranks q' last (set $w' := z'$), or

Structure $p'r'-p'r'-q's'-q's'$: x', y', z', w' are distinct, x', y' rank p' last and r' second-to-last, and z', w' rank q' last and s' second-to-last.

For choosing $\lambda^* \in BC(P)$ and connecting it to μ^* , it will become relevant that among x', y', z' , and y , enough agents obtain a “bad” house. Under Structure $p'-p'-p'$ and $p'q'-p'-q'$, we take $\lambda^* \in PP(P)$ with the first three picking agents being x', y' , and z' (in this order). If y is not equal to x', y' , or z' , we let her choose as fourth. Lastly, let x_p denote the agent that will receive house p in λ^* . Then, summarized, under these two structures, λ^* satisfies:

- $\lambda^* \in BC(P)$.
- $\lambda^*(x') = p'$.
- y obtains at best her fourth-least favorite house from H .
- Under λ^* , at least one agent in $\{x', y', z'\} \setminus \{x, x_p\}$ obtains their third-least favorite house from H at best (due to x', y', z' all picking first).

Under Structure $p'r'-p'r'-q's'-q's'$, we have the same goal but require an in-between step. We first want to use the inherent symmetry to relabel the agents, and then choose some $\lambda \in PP(P)$ as follows:

- If $x \in \{x', y', z', w'\}$, then ensure $x = w'$ by renaming agents.
- If $y \in \{x', y', z', w'\}$, then ensure $y \in \{x', z'\}$ by renaming agents.
- $\lambda(z') = s'$.
- $\lambda(y') = p'$.
- $\lambda(w') = q'$.
- If $y \notin \{x', y', z', w'\}$, then y obtains her fourth-least favorite house at best under λ .

Considering x , we can clearly relabel the sets $\{x', y'\}$ and $\{z', w'\}$ such that $x \notin \{x', y'\}$. If $x \in \{z', w'\}$, then by symmetry of z' and w' , we can relabel agents such that $x \neq z'$. Considering y , we are interested in the situation $y \in \{x', y', z', w'\}$. The first case is $y \in \{x', y'\}$. Here, by symmetry of x' and y' we relabel such that $y = x'$. The second case is $y \in \{z', w'\}$. Here, if $x = w'$, then $y = z'$. Otherwise, $x \notin \{x', y', z', w'\}$. Then, by symmetry of z' and w' , we can relabel such that $y = z'$. This concludes the relabeling of x', y', z' , and w' . If $y \in \{x', y', z', w'\}$, we choose $\lambda \in PP(P)$ via the picking order such that y' obtains p' , w' obtains q' , z' obtains s' , and then x' obtains r' . This order guarantees that $\lambda(w') = q'$, $\lambda(y') = p'$ and $\lambda(z') = s'$. Otherwise, $y \notin \{x', y', z', w'\}$. If y has p' , q' , r' , and s' as her four least favorite houses, we sacrifice x' and let y pick before her: y' gets p' , w' gets q' , z' gets s' , y gets r' , and x' gets whatever. Otherwise, y has one house not in $\{p', q', r', s'\}$ among her four least favorite houses. Then, y can pick as fifth and obtain her fourth least favorite house at best. Also, $\lambda(y') = p'$, $\lambda(w') = q'$, and $\lambda(z') = s'$ is guaranteed. This proves that we can choose λ as desired in all cases.

We now “rotate” λ over x', y', w' and define $\lambda^*(x') = p'$, $\lambda^*(y') = q' \neq y'_{n-1}$, $\lambda^*(w') = \lambda(x')$, and $\lambda^*(v) = \lambda(v)$ for all $v \in N \setminus \{x', y', w'\}$ (including $\lambda^*(z') = s'$). By our invariant, if $y \notin \{x', y', z', w'\}$, $\lambda^*(y) = \lambda(y)$ and she still obtains her fourth-least favorite house at best. Otherwise, $y \in \{x', z'\}$. Since $\lambda^*(x') = p'$, $\lambda^*(z') = s'$, y obtains her second-least favorite house at best. To prove that $\lambda^* \preceq^* \lambda$, take any $N' \supseteq \{x', y', z', w'\}$ of size 5 and let $H' = \lambda(N')$. $\lambda_{N', H'}$ does not belong to an isolated bottom cycle due to $x', y', z', w' \in N'$. λ^* does not belong to an isolated top cycle, as we have $\lambda^*(x') = p'$ and p' is x' 's least preferred house. We call Lemma 8, and obtain the desired claim. In total, under Structure $p'r'-p'r'-q's'-q's'$, we again constructed $\lambda^* \in M$ such that:

- $\lambda^* \in BC(P)$.
- $\lambda^*(x') = p'$.
- Under λ^* , y obtains her fourth-least favorite house from H at best.
- Under λ^* , at least one agent in $\{x', y', z'\} \setminus \{x, x_p\}$ obtains their third-least favorite house from H at best (due to our relabeling ensuring $x \neq x'$, z' and both x' and z' obtaining their second-least favorite house at best).
- Additionally, $\lambda^*(y') \neq h_{y'}(n-1)$.

Moving from μ^ to λ^* .* Roughly speaking, the following claim states that, starting from the chosen $\mu^* \in TC(P)$, we can perform arbitrary swaps between all agents not equal to x or x_p .

Claim: Let $\mu \in M$ with $\mu(x) = p$ and, under Structure $pr-pr-qs-qs$, $\mu(y) \neq h_y(2)$. For any $u, v \in H \setminus \{x, x_p\}$ let $\mu_{u \leftrightarrow v} \in M$ such that $\mu = \mu_{u \leftrightarrow v}$, except for $\mu(u) = \mu_{u \leftrightarrow v}(v)$ and $\mu(v) = \mu_{u \leftrightarrow v}(u)$. Then, $\mu \preceq^* \mu_{u \leftrightarrow v}$.

To prove this claim, let $N' \supseteq \{x, y, z, u, v\}$ be of size 5 (v may coincide with y or z . Fill up arbitrarily, then). Let $H' = \mu(N')$ and denote for brevity $\lambda := \mu_{u \leftrightarrow v}$. Note that $\mu(N') = \lambda(N')$. Further, note that $\mu_{N', H'}$ is not part of an isolated top cycle, since by design y obtains her third-best house at best under Structure $pr-pr-qs-qs$. Similarly, since $x \in N'$ and $\lambda(x) = p$ is her top choice (in H and thus also in H'), $\lambda_{N', H'}$ is not part of an isolated bottom cycle. Hence, by Lemma 8, $\mu \preceq^* \lambda$. This concludes the proof of the claim.

Now, start by setting $\mu^0 := \mu^*$ and $i := 0$. Iterate over each $u \in H \setminus \{x, x_p\}$. Check if $\mu^i(u) = \lambda^*(u)$. If yes, there is nothing to do and we proceed with the next u . Otherwise, identify $v \neq u$ with $\mu^i(v) = \lambda^*(u)$. By the proven claim, $\mu^i \preceq^* \mu_{u \leftrightarrow v}^i$. By construction, we still have $\mu_{u \leftrightarrow v}^i(x) = p$, and it coincides with λ^* in more agents than μ^i . Set $\mu^{i+1} := \mu_{u \leftrightarrow v}^i$, then increment $i \leftarrow i+1$. Note that under Structure $pr-pr-qs-qs$, repeatedly applying the claim requires that $\mu^i(y) \neq h_y(2)$ for all $i \leq k$. For this, first note that $\lambda^*(y) \neq h_y(2)$, as it is y 's fourth-least favorite house in H at best and $n \geq 6$. Hence, we can in the first step pick $u = y$ if needed, and obtain $\mu^1(y) = \lambda^*(y)$. Afterwards, it will not reappear in the procedure as some v : Observe that in step $i+1$, the chosen v with $\mu^i(v) = \lambda^*(u)$ satisfies $\mu^i(v) = \lambda^*(u) \neq \lambda^*(v)$. For $i = 1$, clearly $\mu^1(y) = \lambda^*(y)$, and $v \neq y$. Inductively, $\mu^{i+1}(y) = \mu^i(y) = \lambda^*(y)$ for all i , as desired. In at most $n-2$ steps, we arrive at some $\mu^k \in M$ with $\mu^* \preceq^* \mu^k$, $\mu^k(x) = \lambda^*(x)$ for all $x \in H \setminus \{x, x_p\}$.

Finally, we prove $\mu^k \preceq^* \lambda^*$. To this end, consider $N' = \{x, x_p, x', y', z'\}$ and $H' = \mu^k(N')$. We know μ^k coincides with λ^* on $\{x', y', z'\} \setminus \{x, x_p\}$. Hence, $\mu_{N', H'}^k$ is not part of an isolated top cycle, as at least one agent in N' obtains their third-least favorite house at best. Further, $x', y', z' \in N'$ and under structures $p'-p'-p'$ and $p'q'-p'-q'$, there are no isolated bottom cycles. Finally, regarding structure $p'r'-p'r'-q's'-q's'$, recall that we chose λ^* such that $\lambda^*(y') \neq h_{y'}(n-1)$. It follows that $\lambda_{N', H'}^k$ is not part of an isolated bottom cycle. Calling Lemma 8 on N', H' to obtain $\mu^k \preceq^* \lambda^*$ concludes the proof. $\square$

Proposition 5. *Let P be any profile with $n \geq 5$. Then, $|TC(P)| > 2$ and $|BC(P)| \leq 2$ implies $TC(P) = M \setminus BC(P)$.*

Some lemmas before the main proof are required.

Lemma 9. *Let λ^1 be an assignment such that x, y , or z obtains one of her three most preferred houses, say t , under Structure $p-p-p$ or Structure $pq-p-q$. Then, there is λ with $\lambda^1 \preceq^* \lambda$ and $\lambda(x) = p$.*

PROOF. Consider

$$\begin{aligned} x: & p, h_x(2), h_x(3), \dots \\ P = y: & p, h_y(2), h_y(3), \dots \\ z: & p, h_z(2), h_z(3), \dots \end{aligned}$$

$$P' = \begin{array}{lcl} x: & p, & q, \dots \\ y: & p, & h_y(2), \dots \\ z: & q, & h_z(2), \dots \end{array}$$

Our goal is to show that we can keep each structure intact on a subprofile with five agents and five houses.

For P : Add $v \in N$ with $\lambda^1(v) = p$ to x, y, z and fill x, y, z, v up until it is a set N' of size 5. Then, $\lambda^1(N') \supseteq \{p, t\}$.

For P' : Add $v, w \in N$ with $\lambda^1(v) = p$ and $\lambda^1(w) = q$ to x, y, z and fill x, y, z, v, w up until it is a set N' of size 5 if needed. Then, $\lambda^1(N') \supseteq \{p, q, t\}$.

Clearly, the structure remains intact on N', H' , so there is no isolated top cycle on $M_{N', H'}$. Further, λ^1 is not part of an isolated bottom cycle, as some agent obtains t as their third favorite house or better. It follows by Lemma 8 that λ^1 can reach every assignment λ with $\lambda(N') = H'$, in particular one with $\lambda(x) = p$. $\square$

Lemma 10. *Let λ^1 be an assignment such that one of x, y, z, w obtains one of her two most preferred houses, say t , under Structure $pr-pr-qs-qs$. Then, there is λ with $\lambda^1 \succeq^* \lambda$ and $\lambda(x) = p$.*

PROOF. Consider

$$P'' = \begin{array}{lcl} x: & p, & r, \dots \\ y: & p, & r, \dots \\ z: & q, & s, \dots \\ w: & q, & s, \dots \end{array}$$

W.l.o.g., let w obtain t , one of her top two houses. Let $v \in N$ with $\lambda^1(v) = p$. Then, if needed, fill up x, y, z, w, v to some set N' of size 5. Since $\lambda^1(N') \supseteq \{p, t\}$, the four agents x, y, z, w top-rank two houses twice, and there is no isolated top cycle in $M_{N', H'}$. Further, λ^1 is not in an isolated bottom cycle, as an agent obtains her second-best house or better. Hence, define λ by letting x and v swap houses from λ^1 . Lemma 8 yields $\lambda^1 \succeq^* \lambda$. $\square$

PROOF OF PROPOSITION 5. Now, consider any assignment $\lambda^* \notin BC(P)$.

The following is our **Claim**: $\lambda^* \succeq^* \lambda^1$ for some λ^1 with $\lambda^1(x) = p$.

If $\lambda^*(x) = p$, we are done immediately by $\lambda^1 = \lambda^*$. Otherwise, since λ^* is not Pareto-pessimal, there exists a swap cycle of houses such that all agents on this cycle will be worse off after applying it to λ^* . Our overall approach is to swap houses in a way that x receives p (getting better off) while still enough agents are worse off, so that the assignment after the swap (λ^1) is still weakly majority-dominated by λ^* . To this end, we make use of the just mentioned possible swap which is a Pareto-worsening. We now distinguish several cases depending on the structure. Recall that P has one of the following three forms

$$P^1 = \begin{array}{lcl} x: & p, & h_x(2), h_x(3), \dots \\ y: & p, & h_y(2), h_y(3), \dots \\ z: & p, & h_z(2), h_z(3), \dots \end{array}$$

$$P^2 = \begin{array}{lcl} x: & p, & q, \dots \\ y: & p, & h_y(2), \dots \\ z: & q, & h_z(2), \dots \end{array}$$

$$P^3 = \begin{array}{lcl} x: & p, & r, \dots \\ y: & p, & r, \dots \\ z: & q, & s, \dots \\ w: & q, & s, \dots \end{array}$$

For P^1 : If λ^* assigns p to x, y , or z , then we can w.l.o.g. rename the agents such that $\lambda^*(x) = p$. Otherwise, consider the swap cycle: if it is of length 2, at least one agent among x, y, z is untouched by it, say x . Further, one of x 's top three houses t is not affected by it, and since two agents are worse off by the cycle we can find the agent w with $\lambda^*(w) = t$ and create λ^1 by letting x and w swap their houses simultaneously with the two agents getting worse houses by the cycle. Then, as at most 2 agents improve by letting x and w swap their houses (and they are disjoint from the two agents being worse off), we have $\lambda^* \succeq \lambda^1$. Using Lemma 9, we obtain $\lambda^1 \succeq^* \lambda^{1.5}$ with $\lambda^{1.5}$ being some assignment satisfying $\lambda^{1.5}(x) = p$. If the swap cycle is of length at least 3, it again can either not contain p or none of the three agents x, y, z . If it doesn't contain p , simply let w be the agent such that $\lambda^*(w) = p$, and force swap the houses of w and x after performing the Pareto-worsening swap cycle, to obtain λ^1 . Agent x improves, and at most one agent from the cycle improves. Hence, at least two agents from the cycle are still worse off, yielding $\lambda^* \succeq \lambda^1$. Similarly, if it does not contain w.l.o.g. x , force swap p to x with the same argument.

For P^2 : We can clearly force swap a top three house t to one of the three agents x, y, z : If the Pareto-worsening swap cycle is of length 2, we can use the pigeon hole principle because one of the three agents, w.l.o.g. x , and one of her top three houses t are not affected by the swap cycle. Define λ^1 by identifying the agent currently holding house t and letting her swap houses with x after performing the Pareto-worsening swap cycle. This yields $\lambda^* \succeq \lambda^1$. Then, apply Lemma 9 to obtain $\lambda^{1.5}$, in which $\lambda^{1.5}(x) = p$. If the Pareto-worsening swap cycle is of length 3, it either does not involve one of x, y, z or one of p, q . If, e.g., the cycle does not contain x , we define λ^1 by letting x swap her house with the agent who holds p after performing the Pareto-worsening swap cycle. If the Pareto-worsening swap cycle is of length at least 4, just give agent x house p after performing the Pareto-worsening swap cycle.

For P^3 Again, we perform a case analysis based on the length of the Pareto-worsening swap cycle.

- Consider first the case where the cycle has length 2.
 - * It could be that houses p, r and agents z, w form this swap cycle. This is the most annoying case, as our improving swap can neither contain p or r , nor can we directly swap s to z or w . Effectively, this prevents all four agents x, y, z, w from obtaining one of their top 2 houses directly by means of executing a swap directly after the worsening swap. We are forced to take a detour by defining λ^1 from λ^* by first applying the worsening swap cycle and then letting x obtain her *third*-best house by swapping with some other

agent (which cannot be z, w , as they hold p, r). λ^1 is well defined, and $\lambda^* \succ \lambda^1$. Annoyingly, λ^1 is too weak to call Lemma 10. We hence resort to Lemma 8 to swap p to x or y : Let v be the agent with $\lambda^1(v) = s$. Considering $N' = \{x, y, z, w, v\}$, we see that $H' = \lambda^1(N')$ contains s . Further, by assumption on the Pareto-worsening swap cycle, z, w hold p and r ! Hence, the on $P_{N', H'}$, the structure with two pairs of agents sharing the same top choice is intact, meaning $M_{N', H'}$ contains no isolated top cycle. As λ^1 gives agent x his third ranked house, it is not in an isolated bottom cycle in $M_{N', H'}$. Using Lemma 8, it reaches all assignments with $\lambda(N') = H'$ on the majority graph, in particular $\lambda^{1.5}$ which assigns p to x .

- * If one of p or r is not affected by the Pareto worsening swap cycle, we proceed by a further case analysis. The first case is that the cycle does not affect at least one agent of x and y . The second case is that the cycle is between x and y . In the first case, force swap p or r (whichever one of them is not affected by the Pareto-worsening swap cycle) to x or y (whichever one of them is not affected by the Pareto-worsening swap cycle) after performing the Pareto-worsening swap cycle, to obtain λ^1 . In the second case, we can assume that one of q or s is unaffected by the cycle because this situation would be identical to the case where the cycle lets w and z swap houses p and r up to renaming. We hence force swap q or s to z after performing the Pareto-worsening swap cycle to obtain λ^1 . We make use of Lemma 10, which yields $\lambda^1 \succ^* \lambda^{1.5}$ for some $\lambda^{1.5}$ with $\lambda^{1.5}(x) = p$.
- * Finally, consider the case when at least of the agents w or z is not affected by the Pareto-worsening swap cycle, say w.l.o.g. w . We know that, since we are not in the previous case, the two houses being swapped by the cycle are p and r . This means we can simultaneously perform two swaps: Let w swap with the agent who gets q in λ , and let the two agents getting p and r swap their houses. The result of these two simultaneous swaps is λ^1 and we have $\lambda^* \succ \lambda^1$. By Lemma 10, we can reach a $\lambda^{1.5}$, giving x to p .
- If the cycle is of length 3: Either some agent among x, y, z, w already obtained one of her top 2 choices or there is an easy force swap to achieve this. Call again Lemma 10.
- If the cycle is of length 4 or larger: Force swap p to x after performing the Pareto-worsening swap cycle.

In all cases, we obtain an assignment, overwrite λ^1 , with $\lambda^1(x) = p$. This proves the claim.

With this claim, the proof becomes significantly easier. Take any Pareto-optimal assignment μ^* with $\mu^*(x) = p$ and, if applicable,

$\mu^*(z) = q$. We again perform a case analysis based on the profile structure.

- For P^1 , we start from λ^1 , and always consider x, y, z . We add some w on which λ^i does not coincide with μ^* , and the v which holds $\mu^*(w)$ under λ^i . Force swap to obtain λ^{i+1} , as λ^i is not in an isolated bottom cycle and by $p \in H'$ and $\lambda^1(x) = p$ there is no isolated top cycle.
- For P^2 , we first start by swapping q to z (with the same connectedness argument) to obtain λ^2 . From that point on, any $N' \supseteq \{x, y, z\}$ contains the structure pq - p - q with $H' = \lambda^i N'$. We keep adding w, v to λ^i as in the item before.
- For P^3 , we can swap q to z in the first step (with the same connectedness argument). This case is rather tricky as we need all five agents x, y, z, w, v with $\lambda^1(v) = q$. From this point on, the double double top is intact whenever x and z are present in λ^i . Since x, y, z, w do not allow for the addition of two new agents u, v , we need to swap in two steps. First, we identify some v for which λ^i assigns a different house than μ^* , and find the agent u holding it. Then, we let u and y swap. Then, we let y and v swap. We can do this until all agents $v \neq x, y, z, w$ except for one last u are satisfied. The last five agents x, y, z, w, u we easily deal with at once.

This concludes the proof, as we have shown that from any assignment not in an isolated bottom cycle, we can reach a serial dictatorship on the majority graph. $\square$

C PATH FOR EXAMPLE 4

Here, we provide a concrete path from the assignment μ in Example 4 to a serial dictatorship λ in the majority graph. What follows, is a sequence of pairs of assignments with the red assignment weakly dominating the other one (which is marked in blue for the agents receiving a different house). The profile P of course remains the same. In the first illustration, the red assignment is μ . In the last illustration, the blue assignment is λ .

1: f, b, d, e, c, a, g
2: d, f, g, a, b, e, c
3: d, a, c, g, e, b, f
 $P =$ 4: a, d, b, f, g, e, c
5: c, g, e, b, f, d, a
6: f, a, e, d, g, c, b
7: c, d, e, b, g, f, a

1: f, b, d, e, c, a, g
2: d, f, g, a, b, e, c
3: d, a, c, g, e, b, f
 $P =$ 4: a, d, b, f, g, e, c
5: c, g, e, b, f, d, a
6: f, a, e, d, g, c, b
7: c, d, e, b, g, f, a

1: f , b , d , e , c , a , g
 2: d , f , g , a , b , e , c
 3: d , a , c , g , e , b , f
 $P =$ 4: a , d , b , f , g , e , c
 5: c , g , e , b , f , d , a
 6: f , a , e , d , g , c , b
 7: c , d , e , b , g , f , a

 1: f , b , d , e , c , a , g
 2: d , f , g , a , b , e , c
 3: d , a , c , g , e , b , f
 $P =$ 4: a , d , b , f , g , e , c
 5: c , g , e , b , f , d , a
 6: f , a , e , d , g , c , b
 7: c , d , e , b , g , f , a

1: f , b , d , e , c , a , g
 2: d , f , g , a , b , e , c
 3: d , a , c , g , e , b , f
 $P =$ 4: a , d , b , f , g , e , c
 5: c , g , e , b , f , d , a
 6: f , a , e , d , g , c , b
 7: c , d , e , b , g , f , a

The assignment λ is a serial dictatorship, e.g., with the picking order 3, 4, 6, 7, 1, 5, 2.