

Picking Sequences and Monotonicity in Weighted Fair Division

Mithun Chakraborty¹, Ulrike Schmidt-Kraepelin², and
Warut Suksompong³

¹ University of Michigan, Ann Arbor, USA
`dcsmc@umich.edu`

² Technische Universität Berlin, Berlin, Germany
`u.schmidt-kraepelin@tu-berlin.de`

³ National University of Singapore, Singapore
`warut@comp.nus.edu.sg`

Abstract. We study the problem of fairly allocating indivisible items to agents with different entitlements, which captures, for example, the distribution of ministries among political parties in a coalition government. Our focus is on picking sequences derived from common apportionment methods, including five traditional divisor methods and the quota method. We paint a complete picture of these methods in relation to known envy-freeness and proportionality relaxations for indivisible items as well as monotonicity properties with respect to the resource, population, and weights. In addition, we provide characterizations of picking sequences satisfying each of the fairness notions, and show that the well-studied maximum Nash welfare solution fails resource- and population-monotonicity even in the unweighted setting. Our results serve as an argument in favor of using picking sequences in weighted fair division problems.

Keywords: Picking sequences · apportionment · monotonicity · weighted fair division

1 Introduction

After a national election, the parties forming a coalition government are faced with the task of dividing the ministries among themselves. How can they perform this task in a fair manner, taking into account both their preferences on ministries and the votes that they received in the election?

The study of fairly allocating resources to interested agents (in this case, parties), commonly known as *fair division*, has a long history dating back several decades [13, 26]. Among the most prominent fairness criteria are *envy-freeness*—no agent prefers another agent’s allocated bundle over her own—and *proportionality*—if there are n agents, then every agent receives at least $1/n$ of her value for the entire resource. These criteria implicitly assume that all agents have the same entitlement to the resource, an assumption that is made in the vast majority of the fair division literature, yet utterly fails in our ministry example

as well as when allotting supplies to districts, organizations, or university departments, which typically have different sizes. Fortunately, both envy-freeness and proportionality allow for taking the entitlements, or *weights*, into account in a natural way. For instance, if agent A’s weight is twice that of agent B, then A will be satisfied with respect to weighted envy-freeness as long as she derives at least twice as much value for her own bundle as for B’s bundle. While such weight-based extensions of classical fairness concepts are appropriate for scenarios with different entitlements, they sometimes cannot be satisfied when allocating indivisible items like ministries (e.g., when every party places all of its value on the same important ministry). Consequently, recent work has proposed relaxations including *weighted envy-freeness up to one item (WEF1)* [19] and *weighted proportionality up to one item (WPROP1)* [3], each of which can always be fulfilled.

An attractive class of procedures for allocating items is the class of *picking sequences*: these procedures let agents take turns picking their favorite items according to a prespecified order. Picking sequences are intuitive, can be implemented efficiently, and help preserve privacy since each agent only has to reveal the picks in her turns as opposed to her full preferences. In fact, several methods for apportioning seats in a parliament—a setting commonly known as *apportionment*—can be formulated as picking sequences.⁴ For example, Adams’ method assigns each pick to an agent i who minimizes t_i/w_i , where t_i and w_i denote the number of times that agent i has picked so far and her weight, respectively. Brams and Kaplan [12] proposed using picking sequences to allocate ministries, noting that such sequences have been used in Northern Ireland and Denmark, and Chakraborty et al. [19] showed that the allocation produced by Adams’ method always fulfills WEF1 but not necessarily WPROP1. It is therefore an important question which fairness criteria, if any, are satisfied by picking sequences based on other prevalent apportionment methods.

In addition to fairness, another desirable set of properties for allocation procedures is monotonicity in terms of the parameters of the setting. In particular, *resource-monotonicity* means that whenever an extra item is added, no agent receives a lower utility as a result. Similarly, *population-monotonicity* stipulates that introducing an additional agent should not increase the utility of any existing agent, and *weight-monotonicity* implies that when the weight of an agent increases, her utility does not go down.⁵ Segal-Halevi and Sziklai [29] showed that for divisible items in the unweighted setting, the *maximum Nash welfare (MNW)* solution, which chooses an allocation maximizing the product of the agents’ utilities, is resource- and population-monotone. How do picking sequences and (a

⁴ Note that apportionment is a special case of our setting where all items are identical [7].

⁵ Resource-monotonicity is known as *house-monotonicity* in the context of apportionment; a violation of it is referred to as the *Alabama paradox* [7]. Likewise, violations of (variants of) population- and weight-monotonicity are called the *new states paradox* and the *population paradox*, respectively.

weighted generalization of) MNW perform with respect to monotonicity properties in the weighted allocation of indivisible items?

1.1 Our Results

In this paper, we conduct a thorough investigation of picking sequences based on common apportionment methods, as well as the maximum (weighted) Nash welfare solution, in relation to fairness and monotonicity properties. In addition to WEF1 and WPROP1, we consider *weak weighted envy-freeness up to one item (WWEF1)*, a weakening of WEF1 proposed by Chakraborty et al. [19]. For brevity, we say that an allocation rule satisfies a fairness notion if the allocation that it produces for an arbitrary input instance satisfies that notion.

We begin in Section 3 by establishing fundamental results on our properties in the context of picking sequences. In particular, we define three consistency properties with respect to the resource, population, and weights—for example, resource-consistency means that whenever an item is added, the new picking sequence should simply be the old one with an additional pick appended at the end. We show that resource- and population-consistency imply the respective monotonicity properties for any number of agents, while weight-consistency implies weight-monotonicity only for two agents. In addition, for each fairness notion, we characterize the picking sequences whose output always satisfies that notion.

With this groundwork laid, we proceed to determine the properties satisfied by each allocation rule in Section 4. First, we consider the picking sequences derived from five traditional divisor methods due to Adams, Jefferson, Webster, Hill, and Dean. These methods assign each pick to an agent i who minimizes the ratio $f(t_i)/w_i$ capturing the proportion between the number of times the agent has picked so far and the agent’s weight, where the function f varies from method to method. We establish that all five methods satisfy resource- and population-monotonicity for any number of agents as well as weight-monotonicity for two agents; however, they all fail weight-monotonicity when there are three agents. On the fairness front, all of the methods satisfy WWEF1, but Adams’ is the only one satisfying the stronger notion of WEF1 while Jefferson’s is the only one fulfilling WPROP1.

Next, in Section 5, we address the picking sequence derived from another important apportionment method: the *quota method*. While not itself a divisor method, the quota method has a definition similar to that of Jefferson’s method, which uses the function $f(t) = t + 1$, but also imposes a “quota” to determine each agent’s eligibility. We show that the quota method exhibits similar monotonicity behavior as the divisor methods, with the notable exception that it fails population-monotonicity. As for fairness, like Jefferson’s method, the quota method satisfies WWEF1 and WPROP1. In fact, these two rules are the first to have been shown to satisfy both WWEF1 and WPROP1, to the best of our knowledge.

Finally, in Section 6, we examine the *maximum weighted Nash welfare (MWNW)* solution, which is a natural generalization of the well-studied MNW solution to

the weighted setting. Chakraborty et al. [19] already proved that MWNW satisfies WWEF1 but not WEF1; we show that it fails WPROP1. We then present examples demonstrating that even in the unweighted setting (where MWNW reduces to MNW), the rule fails both resource- and population-monotonicity. This result stands in stark contrast to the aforementioned result of Segal-Halevi and Sziklai [29] that MNW is resource- and population-monotone in the context of divisible items, and is perhaps even more striking given that MNW is known to fulfill several desirable properties [18, 22]. On the positive side, MWNW satisfies weight-monotonicity for any number of agents, and is the only rule to do so among the ones we consider in this paper.

Our results are summarized in Table 1. Overall, we believe that they serve as an argument in favor of using picking sequences in division problems with unequal entitlements in view of both fairness and monotonicity considerations.

	Resource-mon.	Population-mon.	Weight-mon.	WEF1	WWEF1	WPROP1
Adams	✓	✓	✗	✓	✓	✗
Jefferson	✓	✓	✗	✗	✓	✓
Webster	✓	✓	✗	✗	✓	✗
Hill	✓	✓	✗	✗	✓	✗
Dean	✓	✓	✗	✗	✓	✗
Quota	✓	✗	✗	✗	✓	✓
MWNW	✗	✗	✓	✗	✓	✗

Table 1. Summary of our results. Chakraborty et al. [19] showed that Adams’ method satisfies WEF1 and WWEF1 but not WPROP1, while MWNW satisfies WWEF1 but not WEF1. All other results are new to this paper. All rules satisfy weight-monotonicity in the case of two agents. MWNW fails resource- and population-monotonicity even in the unweighted setting.

1.2 Related Work

The fair allocation of indivisible items has received substantial recent attention, notably among computer scientists—see the surveys of Bouveret et al. [9] and Markakis [25]. A large majority of work assumes that all agents have equal entitlements, in which case the notions *envy-freeness up to one item (EF1)* [17, 24] and *proportionality up to one item (PROP1)* [1, 20] are often considered. Both WEF1 and WWEF1 reduce to EF1 in the unweighted setting, while WPROP1 reduces to PROP1. Even though EF1 implies PROP1, Chakraborty et al. [19] showed that no rule can simultaneously satisfy WEF1 and WPROP1. Aziz et al. [3] gave a protocol satisfying WPROP1 along with the economic efficiency notion of *Pareto optimality*, while Babaioff et al. [5] considered competitive equilibrium for agents with different budgets representing their weights. Farhadi et al. [21] proposed a weighted version of *maximin share fairness* [17, 23], and Aziz

et al. [2] studied the analogous notion for chores (i.e., items that yield negative utilities).

Like fair division, apportionment methods have given rise to a long line of work that analyzes their advantages and disadvantages according to various desiderata [7, 27]. As Balinski and Young [7] noted, Adams' method tends to favor agents with smaller weights and Jefferson's typically benefits those with larger weights, whereas the other three divisor methods lie in between. Apportionment has also attracted interest in artificial intelligence [14–16] as well as in philosophy [31].

Finally, picking sequences have been studied by several authors due to their simplicity and practicality [4, 8, 10, 11, 30, 32], with a number of authors investigating manipulation issues. We assume in this paper that agents are not strategic and always pick their most preferred item available. In the unweighted setting, a popular picking sequence is the *round-robin algorithm*, which lets agents pick items in cyclic order until the items run out.

2 Preliminaries

We consider a discrete resource allocation setting with a set of agents $N = [n]$ and a set of indivisible items $M = [m]$, where $[k] := \{1, 2, \dots, k\}$ for any $k \in \mathbb{N}$. Each agent $i \in N$ is endowed with a *weight* $w_i > 0$ and a *utility function* $u_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$; for convenience, we sometimes write $u_i(j)$ instead of $u_i(\{j\})$ for an item $j \in M$. As is very common in the fair division literature, we assume that the utility functions are additive, i.e., $u_i(M') = \sum_{j \in M'} u_i(j)$ for all $i \in N$ and $M' \subseteq M$. An *allocation* $\mathcal{M} = (M_1, \dots, M_n)$ is a partition of the items into n bundles so that agent i receives bundle M_i . An *instance* consists of the agents, items, weights, and utility functions. When all weights are equal (in which case we can take them to be 1 without loss of generality), we refer to the resulting setting as the *unweighted setting*.

We consider the following three fairness notions. The first two notions were proposed by Chakraborty et al. [19] and the third by Aziz et al. [3].

Definition 2.1. *An allocation (M_1, \dots, M_n) is said to satisfy*

- weighted envy-freeness up to one item (WEF1) *if for any $i, j \in N$, there exists $B \subseteq M_j$ with $|B| \leq 1$ such that $\frac{u_i(M_i)}{w_i} \geq \frac{u_i(M_j \setminus B)}{w_j}$;*
- weak weighted envy-freeness up to one item (WWEF1) *if for any $i, j \in N$, there exists $B \subseteq M_j$ with $|B| \leq 1$ such that $\frac{u_i(M_i)}{w_i} \geq \frac{u_i(M_j \setminus B)}{w_j}$ or $\frac{u_i(M_i \cup B)}{w_i} \geq \frac{u_i(M_j)}{w_j}$;*
- weighted proportionality up to one item (WPROP1) *if for any $i \in N$, there exists $B \subseteq M \setminus M_i$ with $|B| \leq 1$ such that $u_i(M_i) \geq \left(\frac{w_i}{\sum_{i' \in N} w_{i'}} \cdot u_i(M) \right) - u_i(B)$.*

Chakraborty et al. [19] showed that no rule can simultaneously satisfy WEF1 and WPROP1. In particular, consider an instance where $m = n$ and every agent

has a nonzero utility for every item. Any WEF1 allocation has to assign exactly one item to each agent.⁶ On the other hand, if a certain agent has the same utility for all items and a sufficiently larger weight than every other agent, WPROP1 will require this agent to receive at least $m - 1$ items.

A *domain* refers to a set of instances. A domain may include all instances with any number of agents and items, weights, and utility functions, or it may only include—for example—all instances with two agents, or all instances with equal weights (this corresponds to the unweighted setting). An *allocation rule* is a function that maps each instance in a given domain to an allocation; it is said to satisfy a fairness notion if the allocation that it produces always fulfills that notion. We now define the three monotonicity properties that we consider—the first two have been studied by Segal-Halevi and Sziklai [28, 29], while the third has not been studied in fair division to the best of our knowledge.

Definition 2.2. *An allocation rule \mathcal{R} with domain \mathcal{I} satisfies*

- resource-monotonicity *if the following holds: for any instance with m items, when an extra item is added as item $m + 1$, if both the original and the modified instance belong to \mathcal{I} , then each agent receives no higher utility from the allocation produced by \mathcal{R} in the original instance than in the modified instance;*
- population-monotonicity *if the following holds: for any instance with n agents, when an extra agent is added as agent $n + 1$, if both the original and the modified instance belong to \mathcal{I} , then each of the first n agents receives at least as much utility from the allocation produced by \mathcal{R} in the original instance as in the modified instance;*
- weight-monotonicity *if the following holds: for any instance, when the weight of an agent increases, if both the original and the modified instance belong to \mathcal{I} , the utility that the agent receives from the allocation produced by \mathcal{R} does not decrease.*

While these monotonicity properties are intuitive and it may seem that any reasonable allocation rule should satisfy them, this is in fact not the case: In the full version of this paper, we show that two popular fair division algorithms—the *envy cycle elimination algorithm* and the *adjusted winner procedure*—fail resource-monotonicity even in the unweighted setting.

Next, we provide definitions related to picking sequences.

Definition 2.3. *A picking sequence on n agents and m items is a sequence $\pi_{n,m,w} = (a_1, a_2, \dots, a_m)$, where $a_i \in N$ for each $i \in M$. A family of picking sequences is a collection $\Pi = \{\pi_{n,m,w}\}$, with at most one picking sequence for each pair of positive integers n, m and each weight vector $\mathbf{w} = (w_1, \dots, w_n)$. A family of picking sequences Π is called*

- resource-consistent *if for every n, m, \mathbf{w} such that both $\pi_{n,m,w}$ and $\pi_{n,m+1,w}$ belong to Π , the sequence $\pi_{n,m,w}$ forms a prefix of $\pi_{n,m+1,w}$;*

⁶ Otherwise an agent with no item will (weighted-)envy an agent with at least two items by more than one item.

- population-consistent if for every $n, m, \mathbf{w}, \mathbf{w}'$ such that $\mathbf{w}' = (w_1, \dots, w_n, w'_{n+1})$ where w'_{n+1} is the weight of agent $n + 1$ and both $\pi_{n,m,\mathbf{w}}$ and $\pi_{n+1,m,\mathbf{w}'}$ belong to Π , the sequence $\pi_{n+1,m,\mathbf{w}'}$ can be obtained from $\pi_{n,m,\mathbf{w}}$ by inserting agent $n + 1$ in some positions (possibly none) and trimming the suffix of the resulting sequence so that the sequence has length m .
- weight-consistent if the following holds: for every n, m, \mathbf{w} and $w'_i > w_i$ such that both $\pi_{n,m,\mathbf{w}}$ and $\pi_{n,m,\mathbf{w}'}$ belong to Π , where $\mathbf{w}' = (w_1, \dots, w'_i, \dots, w_n)$, the sequence $\pi_{n,m,\mathbf{w}'}$ can be obtained from $\pi_{n,m,\mathbf{w}}$ by moving some of agent i 's picks earlier (possibly none), inserting agent i in some positions (possibly none), and trimming the suffix of the resulting sequence so that the sequence has length m .

Given a picking sequence (a_1, \dots, a_m) and the agents' utility functions, we assume that in the i th turn, agent a_i picks her highest-valued item from among the remaining items, breaking ties in a consistent manner (say, in favor of lower-numbered items). We sometimes drop the subscript from $\pi_{n,m,\mathbf{w}}$ when n, m, \mathbf{w} are clear from the context. A family of picking sequences generates an allocation rule, which we will refer to interchangeably with the family itself. We also refer to a picking sequence π interchangeably with the family of picking sequences that consists only of π .

All omitted proofs can be found in the full version of this paper.

3 General Picking Sequences

We begin by proving results for general picking sequences. In addition to being interesting in their own right, these results will later help us determine the properties that each apportionment method satisfies. First, we present characterizations of picking sequences whose output is guaranteed to satisfy each of the fairness notions WEF1, WWEF1, and WPROP1. In particular, we show that a picking sequence guarantees a fairness notion for agents with arbitrary utility functions if and only if it does so for agents with identical utility functions that put utility 1 on some items and 0 on the remaining items. This means that the fairness guarantees for general utilities can be expressed as relatively simple conditions on the number of picks in each prefix of the picking sequence.

Theorem 3.1. *A picking sequence π satisfies WEF1 if and only if for every prefix of π and every pair of agents i, j with $t_j \geq 2$, we have $\frac{t_i}{t_j-1} \geq \frac{w_i}{w_j}$, where t_i and t_j denote the number of agent i 's and agent j 's picks in the prefix, respectively.*

Theorem 3.2. *A picking sequence π satisfies WWEF1 if and only if for every prefix of π and every pair of agents i, j with $t_j \geq 2$, both of the following hold:*

- $\frac{t_i}{t_j-1} \geq \frac{w_i}{w_j}$ if $w_i \geq w_j$;
- $\frac{t_i+1}{t_j} \geq \frac{w_i}{w_j}$ if $w_i \leq w_j$,

where t_i and t_j denote the number of agent i 's and agent j 's picks in the prefix, respectively.

As an example of a picking sequence that satisfies WWEF1 but not WEF1, suppose that $n = 2$, $w_1 = 1$, $w_2 = 2$, and consider the sequence $(1, 2, 2, 2, 2)$. For this sequence, we have $t_1 = 1$ and $t_2 = 4$, and therefore $\frac{t_1+1}{t_2} = \frac{w_1}{w_2} > \frac{t_1}{t_2-1}$.

Theorem 3.3. *A picking sequence π satisfies WPROP1 if and only if for every prefix of π and every agent i , we have $t_i \geq \left(\frac{w_i}{\sum_{i' \in N} w_{i'}} \cdot k \right) - 1$, where t_i and k denote the number of agent i 's picks in the prefix and the length of the prefix, respectively.*

Next, we establish a strong relationship between resource- and population-consistency and the corresponding monotonicity notions.

Theorem 3.4. *Any resource-consistent family of picking sequences satisfies resource-monotonicity.*

Theorem 3.5. *Any population-consistent family of picking sequences satisfies population-monotonicity.*

The relationship between weight-consistency and weight-monotonicity is less straightforward: we show that the former implies the latter in the case of two agents. As we will see later (Proposition 4.2), this relationship breaks down when there are three agents.

Theorem 3.6. *For two agents, any weight-consistent family of picking sequences satisfies weight-monotonicity.*

4 Divisor Methods

As we explained in the introduction, a divisor apportionment method gives rise to a picking sequence that, in each turn, lets an agent i with the smallest $f(t_i)/w_i$ pick the next item (breaking ties in a consistent manner, say, in favor of lower-numbered agents), where t_i denotes the number of times that agent i has picked so far and $f : \mathbb{Z}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a strictly increasing function specific to the method such that $t \leq f(t) \leq t + 1$.⁷ We will refer to the divisor methods and their associated families of picking sequences interchangeably. By definition, it is clear that every divisor method yields a family of picking sequences (for all n, m, \mathbf{w}) that are resource-, population-, and weight-consistent. Theorems 3.4, 3.5, and 3.6 therefore imply the following:

Corollary 4.1. *Every divisor method satisfies resource-monotonicity and population-monotonicity; it also satisfies weight-monotonicity when there are two agents.*

⁷ Some non-divisor apportionment methods such as Hamilton's method do not give rise to a picking sequence and are therefore not useful in our context [12, p. 149].

The five traditional divisor methods of Adams, Jefferson, Webster, Hill, and Dean have the function $f(t)$ equal to t , $t + 1$, $t + \frac{1}{2}$, $\sqrt{t(t+1)}$, and $\frac{t(t+1)}{t+\frac{1}{2}}$, respectively [7, p. 99]. We prove that, perhaps surprisingly, all five methods fail weight-monotonicity in the case of three agents.⁸ This also means that weight-consistency does not imply weight-monotonicity beyond two agents.

Proposition 4.2. *Each of the five traditional divisor methods does not satisfy weight-monotonicity even when there are three agents.*

We now explore how the five divisor methods fare with respect to the three fairness notions, starting with WEF1.

Theorem 4.3. *Of the five traditional divisor methods, Adams' method is the only one satisfying WEF1.*

Next, we show that interestingly, all five methods satisfy WWEF1, meaning that each of them can guarantee fairness beyond the setting with identical items (i.e., apportionment).

Theorem 4.4. *All five traditional divisor methods satisfy WWEF1.*

We now turn to WPROP1, where we illustrate a strong relationship with a notion from the apportionment setting. A picking sequence $\pi_{n,m,\mathbf{w}}$ is said to satisfy *lower quota* if for any $i \in N$, it holds that $t_i \geq \left\lfloor \frac{w_i \cdot m}{\sum_{i' \in N} w_{i'}} \right\rfloor$, where t_i denotes the number of picks in $\pi_{n,m,\mathbf{w}}$ assigned to agent i .

Proposition 4.5. *Let π be a picking sequence such that every prefix of π satisfies lower quota. Then π satisfies WPROP1.*

Since Jefferson's method satisfies lower quota [7, p. 130] and is resource-consistent, any prefix of its associated picking sequence also satisfies lower quota. By Proposition 4.5, the method satisfies WPROP1. We prove that it is the only traditional divisor method to do so.

Theorem 4.6. *Of the five traditional divisor methods, Jefferson's method is the only one satisfying WPROP1.*

5 Quota Method

Although divisor methods are widely used in practice, they do come with an axiomatic downside: no divisor method satisfies an arguably natural axiom known as *quota* [7, p. 130]. A picking sequence satisfies the quota axiom if for every

⁸ Brams and Kaplan [12, p. 157] showed that for $n = 3$, an agent can do worse when her picks move earlier in the picking sequence. However, their example does not correspond to a weight increase with respect to a divisor method and moreover assumes that agents are strategic rather than truthful.

$i \in N$, it holds that $\left\lfloor \frac{w_i \cdot m}{\sum_{i' \in N} w_{i'}} \right\rfloor \leq t_i \leq \left\lceil \frac{w_i \cdot m}{\sum_{i' \in N} w_{i'}} \right\rceil$, where t_i is the number of picks assigned to agent i by the picking sequence—note that the lower bound simply corresponds to the lower quota notion introduced before Proposition 4.5. Motivated by this observation, Balinski and Young [6] proposed the *quota method* which satisfies the quota axiom as well as resource-consistency. Intuitively, this method can be seen as a constrained version of Jefferson’s method where we choose an agent i minimizing $(t_i + 1)/w_i$ over a restricted subset of “eligible” agents. The picking sequence for the quota method is determined iteratively. For each round $k \in [m]$, let t_i be the number of times agent i has picked in rounds $1, \dots, k - 1$. An agent is *eligible* if she would not exceed her upper bound in the quota axiom upon getting an additional pick in round k . Equivalently, the set of eligible agents is $U(\mathbf{w}, \mathbf{t}, k) = \left\{ i \in N \mid t_i < \frac{w_i \cdot k}{\sum_{i' \in N} w_{i'}} \right\}$, where $\mathbf{t} = (t_1, \dots, t_n)$. Among all eligible agents, the next pick is assigned to an agent minimizing $(t_i + 1)/w_i$, breaking ties in a consistent manner. The method trivially satisfies resource-consistency, which by Theorem 3.4 implies the following:

Corollary 5.1. *The quota method satisfies resource-monotonicity.*

However, satisfying the quota axiom comes at a price: in contrast to all divisor methods (Corollary 4.1), the quota method fails population-monotonicity. Moreover, like the five traditional divisor methods (Proposition 4.2), the quota method fails weight-monotonicity for $n = 3$.

Proposition 5.2. *The quota method does not satisfy population-monotonicity. In addition, it does not satisfy weight-monotonicity even when there are three agents.*

As we observed in Section 4, all divisor methods are weight-consistent for any number of agents by definition, and therefore weight-monotone for two agents by Theorem 3.6. In contrast, we show in the full version of this paper that the quota method is not weight-consistent. However, for two agents, we prove that the method is weight-consistent and hence weight-monotone.

Theorem 5.3. *The quota method satisfies weight-consistency and weight-monotonicity when there are two agents.*

Next, we address fairness criteria for the quota method.

Theorem 5.4. *The quota method fails WEF1 but satisfies WWEF1 and WPROP1.*

6 Maximum (Weighted) Nash Welfare

Given any instance in the unweighted setting, the *maximum Nash welfare (MNW)* solution chooses an allocation that maximizes the *Nash welfare*, i.e., the product of the agents’ utilities. MNW is known to satisfy strong fairness guarantees including EF1 [18, 22]. When weights are present, a natural generalization called

maximum weighted Nash welfare (MWNW), which maximizes⁹ the weighted product $\prod_{i=1}^n u_i(M_i)^{w_i}$, satisfies WWEF1 but not WEF1 [19].

Segal-Halevi and Sziklai [29] showed that for *divisible* items in the unweighted setting, MNW satisfies both resource- and population-monotonicity. It is therefore rather surprising that the same is not true for indivisible items.

Proposition 6.1. *In the unweighted setting, MNW satisfies neither resource-monotonicity nor population-monotonicity.*

Proof. For resource-monotonicity, consider four items and two agents, both with weight 1, with the following utilities:

	Item 1	Item 2	Item 3	Item 4
Agent 1	3	2	2	2
Agent 2	2	2	1	1

With only the first three items available, the unique MNW allocation gives items 1 and 3 to agent 1 and item 2 to agent 2, resulting in a utility of 5 for agent 1. However, when we add item 4, MNW uniquely allocates items 3 and 4 to agent 1 and items 1 and 2 to agent 2. So agent 1’s utility drops to 4, violating resource-monotonicity.

The example for population-monotonicity uses four items and three agents and can be found in the full version of this paper.

On the other hand, we prove that MWNW fulfills weight-monotonicity, making it the only rule among the ones we consider in this paper to do so.

Theorem 6.2. *MWNW satisfies weight-monotonicity.*

The proof of Theorem 6.2 is simple: If an allocation \mathcal{M} has a positive weighted Nash welfare, then increasing the weight of agent i from w_i to w'_i changes the weighted Nash welfare of \mathcal{M} by a factor of $u_i(M_i)^{w'_i-w_i}$. The theorem then follows by observing that this quantity is increasing in $u_i(M_i)$.

As mentioned, MWNW is known to satisfy WWEF1 but not WEF1. To complete the picture, we show that it does not satisfy WPROP1. This contrasts with the unweighted setting, where MNW satisfies PROP1 (since EF1 implies PROP1).

Proposition 6.3. *MWNW does not satisfy WPROP1.*

⁹ Ties can be broken arbitrarily unless the maximum weighted Nash welfare is 0 (which occurs, for example, when $m < n$). In this exceptional case, we choose a maximum subset of agents who can be given positive utilities simultaneously, breaking ties in a consistent manner among all such subsets independently of the weights (e.g., lexicographically with respect to the agent indices). We then pick an allocation maximizing the weighted Nash welfare of the agents in this subset.

7 Conclusion and Future Work

In this paper, we have thoroughly investigated picking sequences derived from common apportionment methods, including the five traditional divisor methods and the quota method, in relation to fairness and monotonicity properties. Our results indicate that picking sequences based on divisor methods provide strong guarantees in weighted fair division scenarios such as allocating ministries to political parties, with Adams’ and Jefferson’s methods standing out for fulfilling WEF1 and WPROP1, respectively. Since Jefferson’s method tends to favor large parties while Adams’ often benefits smaller ones (see Section 1.2), an interesting question is whether there are compelling fairness notions in addition to WWEF1 that the other three traditional divisor methods, which intuitively lie somewhere in the middle, satisfy.

A natural direction for future work is to construct rules that exhibit a stronger axiomatic behavior than the ones considered in this paper, or to prove that such rules do not exist. Satisfying the three monotonicity properties simultaneously is trivial: one can always allocate all items to a fixed agent, or ignore the weights and use the round-robin algorithm with a fixed ordering. However, this will not result in a fair allocation with respect to the weights. Does there exist a rule fulfilling the three monotonicity properties along with, say, WWEF1? Other notions that one could consider include strategyproofness and Pareto optimality—even in the unweighted setting, we are not aware of any rule that simultaneously fulfills EF1, Pareto optimality, and resource- or population-monotonicity.¹⁰ Our work leaves many intriguing combinations of properties to investigate, which we hope will lead to more interesting rules for fair resource allocation.

Acknowledgments

This work was partially supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1 and by an NUS Start-up Grant. We would like to thank Dominik Peters for helpful discussion and the anonymous reviewers for valuable feedback.

¹⁰ We remark here that no rule that relies only on agents’ rankings over items, including all picking sequences, can simultaneously satisfy EF1 and Pareto optimality. To see this, consider six items and two agents who rank them in the order $1, 2, \dots, 6$ from best to worst. In order to guarantee EF1, each agent must get one item from each of the sets $\{1, 2\}$, $\{3, 4\}$, $\{5, 6\}$. However, any such allocation is not always Pareto optimal. Indeed, it could be that agent 1 has utilities $100, 99, 3, 2, 1, 0$ for the items, while agent 2 has utilities $100, 99, 98, 97, 96, 95$. Any allocation satisfying the above condition gives agent 1 exactly one of the first two items and agent 2 exactly three items, but this is Pareto dominated by the allocation that gives agent 1 the first two items and agent 2 the last four items.

References

1. Aziz, H., Caragiannis, I., Igarashi, A., Walsh, T.: Fair allocation of indivisible goods and chores. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI). pp. 53–59 (2019)
2. Aziz, H., Chan, H., Li, B.: Weighted maxmin fair share allocation of indivisible chores. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI). pp. 46–52 (2019)
3. Aziz, H., Moulin, H., Sandomirskiy, F.: A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters* **48**(5), 573–578 (2020)
4. Aziz, H., Walsh, T., Xia, L.: Possible and necessary allocations via sequential mechanisms. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI). pp. 468–474 (2015)
5. Babaioff, M., Nisan, N., Talgam-Cohen, I.: Fair allocation through competitive equilibrium from generic incomes. In: Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (FAT*). p. 180 (2019)
6. Balinski, M.L., Young, H.P.: The quota method of apportionment. *American Mathematical Monthly* **82**(7), 701–730 (1975)
7. Balinski, M.L., Young, H.P.: *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press (2001)
8. Beynier, A., Bouveret, S., Lemaître, M., Maudet, N., Rey, S., Shams, P.: Efficiency, sequenceability and deal-optimality in fair division of indivisible goods. In: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 900–908 (2019)
9. Bouveret, S., Chevaleyre, Y., Maudet, N.: Fair allocation of indivisible goods. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) *Handbook of Computational Social Choice*, chap. 12, pp. 284–310. Cambridge University Press (2016)
10. Bouveret, S., Lang, J.: A general elicitation-free protocol for allocating indivisible goods. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI). pp. 73–78 (2011)
11. Bouveret, S., Lang, J.: Manipulating picking sequences. In: Proceedings of the 21st European Conference on Artificial Intelligence (ECAI). pp. 141–146 (2014)
12. Brams, S.J., Kaplan, T.R.: Dividing the indivisible: Procedures for allocating cabinet ministries to political parties in a parliamentary system. *Journal of Theoretical Politics* **16**(2), 143–173 (2004)
13. Brams, S.J., Taylor, A.D.: *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press (1996)
14. Brederick, R., Faliszewski, P., Furdyna, M., Kaczmarczyk, A., Lackner, M.: Strategic campaign management in apportionment elections. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI). pp. 103–109 (2020)
15. Brill, M., Gözl, P., Peters, D., Schmidt-Kraepelin, U., Wilker, K.: Approval-based apportionment. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI). pp. 1854–1861 (2020)
16. Brill, M., Laslier, J.F., Skowron, P.: Multiwinner approval rules as apportionment methods. In: Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI). pp. 414–420 (2017)

17. Budish, E.: The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* **119**(6), 1061–1103 (2011)
18. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation* **7**(3), 12:1–12:32 (2019)
19. Chakraborty, M., Igarashi, A., Suksompong, W., Zick, Y.: Weighted envy-freeness in indivisible item allocation. In: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. pp. 231–239 (2020)
20. Conitzer, V., Freeman, R., Shah, N.: Fair public decision making. In: *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. pp. 629–646 (2017)
21. Farhadi, A., Ghodsi, M., Hajiaghayi, M., Lahaie, S., Pennock, D., Seddighin, M., Seddighin, S., Yami, H.: Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research* **64**, 1–20 (2019)
22. Halpern, D., Procaccia, A.D., Psomas, A., Shah, N.: Fair division with binary valuations: One rule to rule them all. In: *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. pp. 370–383 (2020)
23. Kurokawa, D., Procaccia, A.D., Wang, J.: Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM* **64**(2), 8:1–8:27 (2018)
24. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: *Proceedings of the 5th ACM Conference on Economics and Computation (EC)*. pp. 125–131 (2004)
25. Markakis, E.: Approximation algorithms and hardness results for fair division. In: Endriss, U. (ed.) *Trends in Computational Social Choice*, chap. 12, pp. 231–247. AI Access (2017)
26. Moulin, H.: *Fair Division and Collective Welfare*. MIT Press (2003)
27. Pukelsheim, F.: *Proportional Representation: Apportionment Methods and Their Applications*. Springer (2014)
28. Segal-Halevi, E., Sziklai, B.R.: Resource-monotonicity and population-monotonicity in connected cake-cutting. *Mathematical Social Sciences* **95**, 19–30 (2018)
29. Segal-Halevi, E., Sziklai, B.R.: Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory* **68**(2), 363–401 (2019)
30. Tominaga, Y., Todo, T., Yokoo, M.: Manipulations in two-agent sequential allocation with random sequences. In: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. pp. 141–149 (2016)
31. Wintein, S., Heilmann, C.: Dividing the indivisible: Apportionment and philosophical theories of fairness. *Politics, Philosophy & Economics* **17**(1), 51–74 (2018)
32. Xiao, M., Ling, J.: Algorithms for manipulating sequential allocation. In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*. pp. 2302–2309 (2020)