

The Leximin Approach for a Sequence of Collective Decisions

Ido Kahana^{a,*} and Noam Hazon^{a,**}

^aDepartment of Computer Science, Ariel University, Israel

Abstract. In many situations, several agents need to make a sequence of decisions. For example, a group of workers that needs to decide where their weekly meeting should take place. In such situations, a decision-making mechanism must consider fairness notions. In this paper, we analyze the fairness of three known mechanisms: round-robin, maximum Nash welfare, and leximin. We consider both offline and online settings, and concentrate on the fairness notion of proportionality and its relaxations. Specifically, in the offline setting, we show that the three mechanisms fail to find a proportional or approximate-proportional outcome, even if such an outcome exists. We thus introduce a new fairness property that captures this requirement, and show that a variant of the leximin mechanism satisfies the new fairness property. In the online setting, we show that it is impossible to guarantee proportionality or its relaxations. We thus consider a natural restriction on the agents' preferences, and show that the leximin mechanism guarantees the best possible additive approximation to proportionality and satisfies all the relaxations of proportionality.

1 Introduction

In many situations, a group of agents needs to reach a collective decision, thus needing a collective decision-making mechanism. For example, friends choosing a restaurant may utilize a voting procedure for deciding where to go. Several voting procedures, as well as other collective decision-making mechanisms, have been developed for reaching a one-time, single decision, but there are many situations in which there is a sequence of decisions. For example, consider a city council that needs to decide, each month, which activity to subsidize. It can choose to organize a family activity in the public park, hold a concert in the community center, or operate an overnight bus, but only one activity can be subsidized each month. A senior citizen may benefit the most from a concert, but she may also benefit, to a lesser extent, from the overnight bus. A parent with young kids may benefit the most from an activity in the park, but she may also benefit from a concert. In summary, every activity benefits all the agents, but they may evaluate the activities differently. Clearly, the city council would like to choose activities that satisfy all of the citizens, and it may take advantage of the fact that the decision situation is repeated every month. As another example, consider friends that study together for an exam, and they need to choose a restaurant every day. Clearly, if most of them prefer a pizza, then it is reasonable that they will go together to a pizzeria on many days, even though one of the friends, say Bob, prefers sushi. However, since there is a sequence of

decisions, it is also reasonable to consider fairness, i.e., an outcome in which the group of friends goes to a pizzeria every day may not be fair for Bob.

Arguably, one of the most fundamental notions of fairness is proportionality (*Prop*). That is, assume there is a sequence of decisions, and we call each decision situation a round. There are n agents, and each agent evaluates the candidates in each round. We would like that each one of the agents will get at least a $1/n$ fraction of the utility she would get if she could solely decide on the outcome in each round. Unfortunately, there are instances in which a proportional outcome does not exist, and thus it is reasonable to consider a multiplicative approximation of proportionality, namely, α -*Prop*. Alternatively, Conitzer et al. [11] have suggested two relaxations of proportionality, namely, *Prop1* and *RRS* (see Section 3 for the formal definitions).

When dealing with a sequence of decisions, it is important to distinguish between offline and online settings. In the offline setting, the decision-making mechanism gets as input all the valuations of the agents in all the rounds, and it needs to choose an outcome for every round. In the online setting, at each round, the mechanism gets as an input only the information up to this round, and it thus needs to choose one outcome (for this round). Unfortunately, fairness considerations have been largely neglected in the online setting, and only a few works have analyzed fairness in such a setting. Indeed, it is hard to guarantee fairness in the online setting if the valuations of the agents are unrestricted [6]. One possible restriction to the agents' valuations is based on the Borda scoring rule. Specifically, if there are m candidates, it is assumed that, for each agent, the valuation of her most preferred candidate is $m - 1$, the valuation of her second preferred candidate is $m - 2$, and so on; the valuation of her least preferred candidate is 0. We denote such valuations as Borda valuations. Indeed, the restriction to Borda valuations is also useful in the offline setting, for translating ordinal preference orders to cardinal preferences. That is, there are settings in which it is easier for the agents to express their preferences using ordinal preference orders, in which each agent reports a total order over the set of candidates (i.e., a ranking) for every round. In these settings, several works [26, 5, 13, 12] suggested to translate any ordinal preference order to numerical values according to the Borda scoring rule: the valuation of the highest-ranked candidate is $m - 1$, the valuation of the second-ranked candidate is $m - 2$, and so on.

In this paper, we study collective decision-making mechanisms for a sequence of decisions, both in the offline and online setting, and we analyze three common mechanisms: round-robin (*RR*), maximum Nash welfare (*MNW*), and leximin (*LMin*). We first claim

* Email: ynet.ido@gmail.com

** Email: noamh@ariel.ac.il

that an outcome that satisfies the relaxations of proportionality that were previously suggested (i.e., *Prop1* and *RRS*) might be “far” from being proportional. Specifically, in the offline setting, we show that even though the *RR* and *MNW* mechanisms satisfy *Prop1* and the *LMin* mechanism satisfies *RRS* (with an *RRS*-based normalization), the three mechanisms might fail to find a proportional outcome, even if such an outcome exists. We thus introduce a new natural fairness property, Max-Possible-Prop (*MPP*). According to this property, a mechanism should return an outcome that is as proportional as possible. For example, if an instance admits an outcome that satisfies *Prop*, then the mechanism should return an outcome that satisfies *Prop*. In addition, if an instance does not admit an outcome that satisfies *Prop* but admits an outcome that satisfies $1/2$ -*Prop*, then the mechanism should return an outcome that satisfies $1/2$ -*Prop*. Generally, a mechanism that satisfies *MPP* returns an outcome that is α -proportional, with the maximum possible α for the given instance.

We show that there is an unavoidable trade-off between *MPP* and *Prop1*, and between *MPP* and *RRS* when there are at least 3 agents. That is, a mechanism that satisfies *MPP* does not guarantee any constant factor approximation of *RRS* or *Prop1*. However, we show that a leximin mechanism, in which the valuations are normalized with the proportional value of each agent (i.e., a *Prop*-based normalization), satisfies *MPP*. Moreover, with two agents, this mechanism satisfies *MPP*, *RRS*, and $1/2$ -*Prop1*.

We then analyze the restricted setting of Borda valuations. In this setting, an outcome that satisfies *Prop1* is “far” from being proportional by an additive factor of at most $m - 1$, and we thus consider an additive approximation of *Prop*. We show that, unfortunately, both *RR* and *MNW* do not guarantee an additive constant-factor approximation of *Prop*. On the other hand, the leximin mechanism satisfies *MPP*, *RRS*, *Prop1*, and it guarantees a 1-additive approximation of *Prop*, which is the best possible constant factor additive approximation of *Prop*. We note that Conitzer et al. [11] have raised an open question: is there a mechanism that satisfies *PO*, *Prop1* and *RRS* simultaneously? We partially solve this question- if we restrict the valuations to Borda valuations, then *LMin* is such a mechanism, since we show that it satisfies *PO*, *Prop1*, and *RRS*.

In the online setting, we show that it is impossible to achieve proportionality and even the weaker fairness properties (i.e., *Prop1*, *RRS*, and *MPP*). We thus consider the restriction to Borda valuations, and show that the online leximin mechanism guarantees a 1-additive approximation of *Prop*. Moreover, the online leximin mechanism satisfies *Prop1* and *RRS*. However, we show that there is no online mechanism that satisfies *MPP*, even with Borda valuations.

The main contributions of this paper are threefold. We introduce a natural fairness property, which is (arguably) a better relaxation of proportionality than *RRS* and *Prop1*, and show that a variant of the offline leximin mechanism satisfies it. We also partially solve an open question of Conitzer et al. [11], by using the restriction to Borda valuations. Finally, in the online setting with Borda valuations, we show that the leximin mechanism guarantees the best possible constant factor additive approximation of *Prop*.

Tables 1, and 2 summarize our results. Note that the full proofs of some of the theorems are deferred to the full version of the paper [19].

2 Related Works

The analysis of collective decision-making mechanisms for a sequence of decisions has been studied both in the domain of fair di-

	<i>PO</i>	<i>Prop</i>	<i>Prop1</i>	<i>MPP</i>	<i>RRS</i>
$LMin_{off}^{Prop, n=2}$	✓	✗	$1/2$	✓	✓
$LMin_{off}^{Prop, n>2}$	✓	✗	✗	✓	✗
$LMin_{off}^{RRS}$	✓	✗	$1/2$	✗	✓
RR_{off}	✗	✗	✓	✗	✓
MNW_{off}	✓	✗	✓	✗	$1/n$

Table 1. Summary of results for the offline setting, where there are no restrictions on the valuation. The results in gray are due to [11].

	<i>PO</i>	<i>Prop</i>	<i>Prop1</i>	<i>MPP</i>	<i>RRS</i>
$LMin_{off}$	✓	1-additive	✓	✓	✓
RR_{off}	✗	m -1-additive	✓	✗	✓
MNW_{off}	✓	x -additive $m-1 \geq x$ $x \geq \frac{m-3}{2}$	✓	✗	✗
$LMin_{on}$	✗	1-additive	✓	✗	✓
RR_{on}	✗	m -1-additive	✓	✗	✓
MNW_{on}	✗	x -additive $x \geq \frac{m-3}{2}$?	✗	✗

Table 2. Summary of results with Borda valuations. The results in gray are due to [11].

vision, in which it is commonly called public decision-making [11], and in the domain of voting, in which it is commonly called perpetual voting [21]. Specifically, Conitzer et al. [11] introduce the problem of public decision-making in the offline setting. They propose *RRS* and *Prop1* as relaxations of proportionality, and provide a comprehensive analysis of the *RR*, *MNW*, and *LMin* mechanisms. We note that their analysis of the *LMin* mechanism uses an *RRS*-based normalization, while we propose a *Prop*-based normalization. In addition, we provide an alternative relaxation of proportionality, *MPP*, analyze the restricted setting of Borda valuations, and the online setting. Freeman et al. [16] study the online version of public decision-making. They concentrate on maximizing the Nash welfare, and present two greedy mechanisms. However, they do not analyze their mechanisms with regard to proportionality or its relaxations.

In the domain of voting, Lackner [21] suggests several online voting rules when there is a sequence of decisions, and analyzes them via three axiomatic properties, as well as a quantitative evaluation by computer simulations. In a subsequent paper, Lackner and Maly [22] define two classes of voting rules that are particularly easy to explain to voters, and define specific proportionality axioms. Bulteau et al. [10] study the offline setting, and analyze the fairness for subgroups of voters by adapting the well-established Justified Representation (JR) and Proportional Justified Representation (PJR) axioms. Skowron and Górecki [28] also study the offline setting, and propose a proportionality notion that ensures guarantees to all groups of voters. All of these works assume that the voters express approval preferences, while we study cardinal preferences (or ordinal preferences that are translated to cardinal preferences with the Borda scoring rule).

Our model of public decision-making generalizes the framework

of fair division of private goods, which has been extensively studied [2]. Indeed, fairness considerations have not received much attention in the online setting. Kash, Procaccia, and Shah [20] study online fair division of resources, in which agents arrive at different times. However, unlike in our setting, in their model the agents' preferences do not change over time. Aleksandrov et al. [1] investigate an online fair division problem where items are presented one at a time, and the agents may declare either yes or no for each item. In our setting, there are multiple candidates per round, and agents can express more general utilities. Benade et al. [6] consider indivisible items that arrive online, and assume additive utility, like in our setting. They show that it is impossible to achieve the fairness notion of envy-free up to one item (*EF1*), and design an algorithm that minimizes the maximum envy over time. We strengthen the impossibility result, as we show that with no restriction on the valuations, there is no online mechanism that satisfies even the relaxed fairness properties (i.e., *RRS* or *Prop1*). He et al. [18] show that it is possible to circumvent the impossibility result of [6] by allowing the reallocation of previously allocated items.

Our model is closely related to other frameworks that have been studied in computational social choice. Specifically, the model of voting in combinatorial domains studies the problem of making decisions on several issues [25]. However, the main focus is on the representation of the agents' preferences, since it is assumed that there is a dependency between decisions on different issues. The framework of multi-winner voting [15, 24] considers fairness properties, and the outcome consists of several winning candidates, as in our setting. However, in multi-winner voting, a candidate cannot be usually elected multiple times, as in our setting (but see an exception in [9]). Notably, Bredereck et al. [8] study a sequence of multi-winner elections, in which the difference between the winners in consecutive rounds is upper-bounded. The framework of participatory budgeting [3, 27], which generalizes multi-winner voting, utilizes voting systems for deciding on the funding of public projects. Lackner et al. [23] study a sequence of participatory budgeting problems, and introduce a theory of fairness for this setting. Our setting can also be applied for modeling a sequence of voting on the funding of public projects, but there is no budget constraint. The model of fair allocation of indivisible public goods [14] generalizes participatory budgeting as well as our model of public decision-making. In this model, there is a set of public goods and feasibility constraints on what subsets of goods can be chosen. In the offline setting, Fain et al. [14] provide an additive approximation to the core, which is a fairness notion for groups of agents. Garg et al. [17] consider the simple constraint in which the number of public goods is bounded, and analyze the maximum Nash welfare and leximin objectives with regard to the *RRS* and *Prop1* fairness properties. Banerjee et al. [4] study an online version of fair allocation of public goods, but they analyze fractional outcomes, i.e., the outcome of each round is a mixture of the candidates.

3 Definitions

Consider a set of agents $N = \{1, 2, \dots, n\}$, and a set of candidates $C = \{c_1, c_2, \dots, c_m\}$ ¹. For every round $t = 1, \dots, T$, every agent i reports her valuation $v_i^t(c_j) \in \mathbb{R}_{\geq 0}$ for every candidate c_j . We assume that every agent i has at least one positive valuation (overall). For a given round t , we can write the valuations of all the agents in a

matrix, denoted by V^t , where $V^t = (v_i^t(c_j))_{i,j}$. We denote by $\mathbf{v}^t(c_j)$ the vector of valuations that all the agents assign to candidate c_j . We investigate mechanisms that choose an outcome $\mathbf{o} = (o^1, \dots, o^T)$, $o^t \in C$, which is a choice of a candidate for every round. We assume that the agents have additive utility functions. Therefore, we define the accumulated utility of an agent i from outcome \mathbf{o} , denoted by $u_i(\mathbf{o})$, as $u_i(\mathbf{o}) = \sum_{t=1}^T v_i^t(o^t)$. The accumulated utility vector of all the agents, denoted by $\mathbf{u}(\mathbf{o})$, is a vector in which the i -th entry is $u_i(\mathbf{o})$.

In the offline setting, the mechanism gets as input all of the agents' valuations in all the rounds, i.e., it gets the vector (V^1, \dots, V^T) , and chooses the outcome \mathbf{o} . In the online setting, the mechanism determines the outcome \mathbf{o} sequentially, i.e., the mechanism determines each o^t at round t , since the mechanism does not get the entire input upfront. Indeed, at each round t , the mechanism gets as an input only the information up to this round, i.e., the vector (V^1, \dots, V^t) , and the vector of chosen candidates $\mathbf{o}^{t-1} = (o^1, \dots, o^{t-1})$. We slightly abuse notation and define the accumulated utility of agent i up to round t from outcome \mathbf{o}^{t-1} as $u_i(\mathbf{o}^{t-1}) = \sum_{k=1}^{t-1} v_i^k(o^k)$. Similarly, the accumulated utility vector of all the agents up to round t is $\mathbf{u}(\mathbf{o}^{t-1})$. Clearly, at the first round, the utility of every agent i is zero (i.e., $u_i(\mathbf{o}^0) = 0$).

There are settings in which it is easier for the agents to express their preferences using ordinal preference orders. That is, each agent reports for every round a total order over C . In these settings, we translate the ordinal preference orders to cardinal preferences with the Borda scoring rule. That is, for each agent, the valuation of a candidate c is the number of candidates ranked below c . We denote such valuations as Borda valuations.

4 Efficiency and Fairness

When analyzing offline or online mechanisms, we focus on popular notions of efficiency and fairness. Specifically, for efficiency, we consider the notion of Pareto optimality, which is defined as follows:

Definition 1. An outcome \mathbf{o} is Pareto Optimal (*PO*) if there does not exist another outcome \mathbf{o}' such that for every agent i , (1) $u_i(\mathbf{o}') \geq u_i(\mathbf{o})$ and (2) there exists an agent j such that $u_j(\mathbf{o}') > u_j(\mathbf{o})$.

A mechanism satisfies *PO* if it always chooses an outcome that is *PO*. Generally, we say that a mechanism satisfies a given efficiency or fairness property if it always chooses an outcome that satisfies this property.

For fairness, we concentrate on proportionality, which requires that each agent will receive at least $\frac{1}{n}$ fraction of the utility she would receive if she could choose the outcome. Given an agent i , let $cMax_i^t$ be a candidate with the highest valuation at round t , i.e., $cMax_i^t \in \arg \max_{c \in C} v_i^t(c)$.

Definition 2. Let $Prop_i = \frac{1}{n} \sum_{t=1}^T v_i^t(cMax_i^t)$. For $\alpha \in (0, 1]$, we say that an outcome \mathbf{o} satisfies α -proportionality (α -*Prop*) if for every agent i , $u_i(\mathbf{o}) \geq \alpha \cdot Prop_i$.

We denote 1-*Prop* by *Prop*. Unfortunately, there are instances in which there is no outcome that satisfies α -*Prop*, for any constant α (i.e., an α that does not depend on the given instance). Therefore, Conitzer et al. [11] propose a relaxation of α -*Prop*, which is proportionality up to one round.

Definition 3. An outcome \mathbf{o} satisfies α -*Prop1* if for every agent $i \in N$ there exists a round t such that changing o^t to $cMax_i^t$ ensures that agent i receives a utility of at least $\alpha \cdot Prop_i$, i.e., $\forall i \in N \exists t$ such that $u_i(\mathbf{o}) - v_i^t(o^t) + v_i^t(cMax_i^t) \geq \alpha \cdot Prop_i$.

¹ For the clarity of presentation, we assume that the set of candidates is static. All of our results are easily adapted to the setting in which the set of candidates changes from round to round.

We denote 1-Prop1 by *Prop1*. Indeed, a mechanism may satisfy α -Prop1, but it may still choose an outcome that does not satisfy α -Prop, even though such an outcome exists. We thus propose a new fairness property, Max-Possible-Prop (*MPP*).

Definition 4. An outcome \mathbf{o} satisfies Max-Possible-Prop (*MPP*) if for each α with which there is an outcome that satisfies α -Prop, \mathbf{o} also satisfies α -Prop.

That is, an outcome that satisfies the *MPP* property satisfies α -Prop with the maximum possible α for the given instance.

We also propose to consider an additive approximation of *Prop*. Formally, an outcome \mathbf{o} satisfies a β -additive approximation of *Prop* if for every agent i , $u_i(\mathbf{o}) + \beta \geq \text{Prop}_i$, for $\beta > 0$.

Another notion of fairness is round-robin share (*RRS*) [11]. For an agent i , let \mathbf{cMax}_i be a vector that contains all the values $v_i^t(\mathbf{cMax}_i^t)$, $1 \leq t \leq T$, sorted in a non-ascending order. We denote by $\mathbf{cMax}_i(t)$ the element in the t -th entry of \mathbf{cMax}_i .

Definition 5. Let $\text{RRS}_i = \sum_{1 \leq t \leq \lfloor T/n \rfloor} \mathbf{cMax}_i(t \cdot n)$. For $\alpha \in (0, 1]$ we say that an outcome \mathbf{o} satisfies α -round-robin share (α -RRS) if for every agent i , $u_i(\mathbf{o}) \geq \alpha \cdot \text{RRS}_i$.

We denote 1-RRS by *RRS*.

5 Mechanisms

We concentrate on three families of mechanisms, where each family consists of offline and online mechanisms.

Round-robin (RR) In this mechanism there is a given order over the agents, and they take turns according to this order. In the offline RR, when an agent's turn arrives she chooses a round (that has not been chosen yet), and determines the winning outcome for this round. The agent chooses a round that yields her the highest utility [11]. The online RR chooses a single outcome at each round t , and the order over the agents associates the rounds with the agents. Thus, when an agent's turn arrives, she only determines the winning outcome for the associated round t . Formally, let π be a permutation over N . The offline RR, denoted by RR_{off} , chooses an outcome $\mathbf{o} = (o^1, \dots, o^T)$ according to Algorithm 1.

Algorithm 1 Round-robin offline

```

1:  $\text{rounds} \leftarrow \{1, \dots, T\}$ 
2:  $k \leftarrow 1$ 
3: while  $|\text{rounds}| > 0$  do
4:    $i \leftarrow \pi(1 + (k - 1) \pmod{n})$ 
5:    $t \leftarrow \arg \max_{r \in \text{rounds}} v_i^r(\mathbf{cMax}_i^r)$ 
6:    $\mathbf{o}^t \leftarrow \mathbf{cMax}_i^t$ 
7:    $\text{rounds} \leftarrow \text{rounds} \setminus \{t\}$ 
8:    $k \leftarrow k + 1$ 
9: end while
```

Given a permutation π , let $\text{turn}^t(\pi)$ be a function that associates a round t with an agent, according to the order π , $\text{turn}^t(\pi) = \pi(1 + (t - 1) \pmod{n})$. The online RR is defined as follows.

Definition 6. Given a permutation π over N , the online RR (RR_{on}) chooses an outcome \mathbf{o}^t for round t such that $\mathbf{o}^t = \mathbf{cMax}_{\text{turn}^t(\pi)}^t$.

Maximum Nash welfare (MNW) The Nash welfare of an outcome is the product of the utilities of all the agents. The offline MNW mechanism (MNW_{off}) chooses an outcome that maximizes the Nash welfare. The online MNW mechanism (MNW_{on}) chooses an outcome \mathbf{o}^t for round t that maximizes the Nash welfare up to round t . Formally,

Definition 7. MNW_{off} chooses an outcome \mathbf{o} such that $\mathbf{o} \in \arg \max_{\mathbf{o}'} \prod_{i \in N} u_i(\mathbf{o}')$.

Definition 8. MNW_{on} chooses an outcome \mathbf{o}^t for round t such that $\mathbf{o}^t \in \arg \max_{\mathbf{c}_j \in C} \prod_{i \in N} (u_i(\mathbf{o}^{t-1}) + v_i^t(\mathbf{c}_j))$.

When all outcomes have a Nash welfare of 0, MNW_{off} and MNW_{on} find the largest set of agents that there is an outcome that gives them positive utilities, and choose an outcome that maximizes the product of the agents' utilities. Note that both RR and MNW mechanisms are scale-free. That is, the units of measurement used by the agents for expressing their valuations do not affect the outcome.

Leximin Generally, the motivation behind the mechanisms in this family is to maximize the utility of the agent that has the minimum utility, i.e., the worst-off agent. However, since there might be several such outcomes, a leximin mechanism chooses an outcome that also maximizes the utility of the second worst-off agent, and then the third, and so forth. This idea is formalized by the leximin ordering. Given a vector \mathbf{u} , let $\tilde{\mathbf{u}}$ be the vector \mathbf{u} when the elements are ordered in a non-descending order.

Definition 9. The leximin ordering, \succ , is a total preorder, in which for any two vectors with the same number of elements, \mathbf{x}, \mathbf{y} , $\mathbf{x} \succ \mathbf{y}$ if there exists an index i such that $\tilde{x}_i > \tilde{y}_i$ and $\tilde{x}_j = \tilde{y}_j$ for all $j < i$.

The offline leximin mechanism chooses an outcome such that the vector of accumulated utilities is maximal according to the leximin ordering. Formally, let $\arg \max^\succ$ be the maximum elements under the leximin ordering.

Definition 10. The offline leximin, denoted by LM_{off} , chooses an outcome such that $\mathbf{o} \in \arg \max_{\mathbf{o}'}^\succ u(\mathbf{o}')$.

The offline leximin mechanism is not scale-free, and thus a very high (or low) valuation might bias the mechanism. Therefore, we consider two normalization methods, using either the RRS_i or Prop_i values. The normalization using the RRS_i values is suggested by [16], but it does not make the offline leximin mechanism scale-free. We thus consider also a normalization with the Prop_i values, which makes the offline leximin mechanism scale-free. Let $u\text{RRS}(\mathbf{o})$ be the vector of utilities $u(\mathbf{o})$, in which each $u_i(\mathbf{o})$ is divided by RRS_i . Similarly, $u\text{Prop}(\mathbf{o})$ is the vector of utilities in which each $u_i(\mathbf{o})$ is divided by Prop_i .

Definition 11. The offline leximin *RRS*, denoted by $\text{LM}_{\text{off}}^{\text{RRS}}$, chooses an outcome such that $\mathbf{o} \in \arg \max_{\mathbf{o}'}^\succ u\text{RRS}(\mathbf{o}')$.

Definition 12. The offline leximin *Prop*, denoted by $\text{LM}_{\text{off}}^{\text{Prop}}$, chooses an outcome such that $\mathbf{o} \in \arg \max_{\mathbf{o}'}^\succ u\text{Prop}(\mathbf{o}')$.

In the online setting, since the mechanism does not get the entire input upfront, normalizing the input up to a specific round is meaningless. Moreover, since we show that it is impossible to satisfy proportionality or its relaxations with unrestricted valuations (Theorems 12 and 13), we study the online setting only with Borda valuations. Therefore, the online leximin is defined without normalization.

² If $\text{RRS}_i = 0$, we replace it with an infinitesimal quantity ϵ [16].

Definition 13. The online leximin, denoted by $LMin_{on}$, chooses an outcome o^t for round t such that $o^t \in \arg \max_{c_j \in C} (u(o^{t-1}) + v^t(c_j))$.

Note that the outcome of RR_{off} can be computed efficiently. However, it is intractable to compute the outcomes of MNW_{off} and $LMin_{off}^{RRS}$, due to [11]. It is also intractable to compute the outcomes of $LMin_{off}^{Prop}$, since $LMin_{off}^{Prop}$ satisfies MPP (Theorem 5), and it can thus decide whether a proportional outcome exists, which is a computationally hard problem [7]. Clearly, the outcome of all of the online mechanisms can be computed efficiently.

6 Connections between Fairness Properties

Before analyzing the mechanisms, we show the connections between the various fairness properties of a given outcome, as can be shown in Figure 1.




Figure 1. A graph depicting the connections between the fairness properties. x -additive means x -additive approximation of $Prop$. The dashed edges are only valid with Borda valuations. The blue edges are due to [11]. The red edges are proved in the full version of the paper [19].

Specifically, Conitzer et al. [11] show that $Prop$ entails RRS , and RRS entails $1/2$ - $Prop1$. Clearly, an outcome that satisfies a fairness property satisfies all of its relaxations and approximations. We show additional connections when the valuations are Borda valuations, and we begin with MPP . Note that with Borda valuations, $RRS_i = \lfloor \frac{T}{n} \rfloor (m-1)$ and $Prop_i = \frac{T(m-1)}{n}$, for every agent i .

Theorem 1. With Borda valuations, MPP entails 1-additive approximation of $Prop$.

Proof. Suppose that o satisfies MPP . Clearly, if $Prop_i \leq 1$, then o satisfies a 1-additive approximation of $Prop$. If $Prop_i > 1$, let o' be an outcome that satisfies a 1-additive approximation of $Prop$. Note that such an outcome always exists due to Theorem 14. By definition, $u_i(o') + 1 \geq Prop_i$ for any agent i . That is, $u_i(o') \geq \frac{Prop_i-1}{Prop_i} \cdot Prop_i$. Since $Prop_i > 1$, $\frac{Prop_i-1}{Prop_i} > 0$. In addition, $Prop_i = \frac{T(m-1)}{n}$ for any agent i , and we can thus write that o' satisfies $\frac{Prop_i-1}{Prop_i}$ - $Prop$. Now, since o satisfies MPP , it must also satisfy $\frac{Prop_i-1}{Prop_i}$ - $Prop$. Therefore, $u_i(o) \geq \frac{Prop_i-1}{Prop_i} \cdot Prop_i$, which implies that $u_i(o) + 1 \geq Prop_i$. That is, o satisfies a 1-additive approximation of $Prop$. \square

Theorem 2. With Borda valuations, 1-additive approximation of $Prop$ entails $Prop1$.

Proof. Suppose that o satisfies a 1-additive approximation of $Prop$. Clearly, if o satisfies $Prop$, it also satisfies $Prop1$. If o does not

satisfy $Prop$, then for any agent i in which $u_i(o) < Prop_i$ there exists a round t such that $v_i^t(o^t) < v_i^t(cMax_i^t)$. Since we use Borda valuations, $v_i^t(cMax_i^t) - v_i^t(o^t) \geq 1$. In addition, since o satisfies a 1-additive approximation of $Prop$, then $u_i(o) + 1 \geq Prop_i$. Combining the two inequalities, we get that $(u_i(o) + 1) + (v_i^t(cMax_i^t) - v_i^t(o^t)) \geq Prop_i + 1$. That is, o satisfies $Prop1$. \square

Theorem 3. With Borda valuations, $Prop1$ entails $(m-1)$ -additive approximation of $Prop$.

Proof. Suppose that o satisfies $Prop1$. That is, for each agent $i \in N$, there exists a round t such that $u_i(o) - v_i^t(o^t) + v_i^t(cMax_i^t) \geq Prop_i$. With Borda valuations, for any agent i and round t , $v_i^t(cMax_i^t) = m-1$ and $v_i^t(o^t) \geq 0$. Therefore, $u_i(o) + (m-1) \geq Prop_i$. That is, o satisfies $(m-1)$ -additive approximation of $Prop$. \square

7 Offline Setting

We begin by showing that RR_{off} , MNW_{off} , and $LMin_{off}^{RRS}$ might not find a proportional outcome, even if such an outcome exists. Specifically, consider the following examples.

Example 1. Let $m = 6, n = 2, T = 1$, and

$$V^1 = \begin{pmatrix} 5 & 4 & 3 & 2 & 1 & 0 \\ 2 & 1 & 3 & 0 & 4 & 5 \end{pmatrix}.$$

The only outcome that satisfies $Prop$ is $o = (c_3)$. However, MNW_{off} chooses (c_1) , and RR_{off} chooses either (c_1) or (c_6) .

Example 2. Let $m = 3, n = 2, T = 2$, and

$$V^1 = \begin{pmatrix} 32 & 16 & 8 \\ 2 & 4 & 8 \end{pmatrix}, \quad V^2 = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

The only outcome that satisfies $Prop$ is $o = (c_2, c_2)$. However, since $RRS_1 = RRS_2 = 1$, then $LMin_{off}^{RRS}$ chooses the candidate c_3 in the first round.

It is not a coincidence that RR_{off} , MNW_{off} , and $LMin_{off}^{RRS}$ might fail to find a proportional outcome. We claim that, in general, these mechanisms might output an outcome that is “far” from being proportional, as formally captured by the MPP property. Moreover, note that RR_{off} and $LMin_{off}^{RRS}$ satisfy RRS , and MNW_{off} satisfies $Prop1$ [11]. We show unavoidable tradeoffs between MPP and α - RRS , for any constant α , and between MPP and α - $Prop1$, for any constant α .

The proof is based on the intuition that both $Prop1$ and RRS might fail to find a proportional outcome when there is a round t in which an agent i assigns a relatively high valuation to one of the candidates. In this case, $Prop1$ can be satisfied while ignoring the valuations (of agent i) in the other rounds, which might result in an unproportional outcome. RRS_i , in this case, does not depend on the valuation of i in round t (since the rounds are sorted in a non-ascending order), which might also result in an unproportional outcome. However, the MPP fairness property must consider the valuations of all the agents, in all of the rounds.

Theorem 4. If $n > 2$, a mechanism that satisfies MPP does not satisfy α - RRS , for any constant α , and it does not satisfy α - $Prop1$, for any constant α .

Proof sketch. We build a scenario with n agents, and three distinct agents, i, j , and k . In the first round, agents j and k assign very high valuations to different candidates, while in all other rounds they assign very low valuations to all the candidates. Note that it is impossible to satisfy both j and k in the first round, and assume that the candidate favored by j is selected in the first round. A mechanism that satisfies *MPP* must compensate k in the other rounds. As a result, the utility of agent i is reduced, but still in a way that guarantees her some minimal utility, and overall, all three agents receive this minimal utility. In contrast, a mechanism that satisfies *Prop1* or *RRS* ignores the first round due to the nature of these fairness properties and, as a result, agent i ends up with a utility that is not very low, while agent k ends up with a much lower utility. \square

Since we showed that RR_{off} , MNW_{off} , and $LMin_{off}^{RRS}$ do not satisfy *MPP* even with 2 agents, we introduce $LMin_{off}^{Prop}$, which satisfies *MPP*.

Theorem 5. $LMin_{off}^{Prop}$ satisfies *MPP*.

Proof. Suppose that there exists an α with which there is an outcome \mathbf{o}' that satisfies α -*Prop*. That is, the minimal utility value in $uProp(\mathbf{o}')$ is at least α . By definition, the outcome \mathbf{o} that is returned by $LMin_{off}^{Prop}$ satisfies $uProp(\mathbf{o}) \succeq uProp(\mathbf{o}')$. Therefore, the minimal utility value in $uProp(\mathbf{o})$ is at least α . That is, \mathbf{o} satisfies α -*Prop*. Therefore, $LMin_{off}^{Prop}$ satisfies *MPP*. \square

Clearly, $LMin_{off}^{Prop}$ also satisfies Pareto optimality. However, since $LMin_{off}^{Prop}$ satisfies *MPP*, it does not satisfy α -*RRS* or α -*Prop1*, for any constant α . Indeed, the proof of Theorem 4 assumes that the number of agents, n , is at least 3. If there are only two agents, $LMin_{off}^{Prop}$ satisfies *RRS* and $1/2$ -*Prop1*.

Theorem 6. $LMin_{off}^{Prop}$ with 2 agents satisfies *RRS* and $1/2$ -*Prop1*.

Proof. Clearly, in RR_{off} , the first agent according to the turn's order, π , receives a utility that is at least his *Prop* value. That is, if agent 1 is the first agent according to π , then RR_{off} returns an outcome that guarantees a utility of at least $Prop_1$ to agent 1, and a utility of at least RRS_2 to agent 2. Similarly, if agent 2 is the first agent according to π , then RR_{off} returns an outcome that guarantees a utility of at least $Prop_2$ to agent 2, and a utility of at least RRS_1 to agent 1. Therefore, the outcome \mathbf{o} that is returned by $LMin_{off}^{Prop}$ satisfies $uProp(\mathbf{o}) \succeq (1, \frac{RRS_2}{Prop_2})$ and $uProp(\mathbf{o}) \succeq (\frac{RRS_1}{Prop_1}, 1)$. Now, since $Prop_i \geq RRS_i$ [11], and there are exactly two agents, then $uProp(\mathbf{o}) \succeq (\max(\frac{RRS_2}{Prop_2}, \frac{RRS_1}{Prop_1}), \max(\frac{RRS_2}{Prop_2}, \frac{RRS_1}{Prop_1}))$. Hence, $uProp_i(\mathbf{o}) \geq \max(\frac{RRS_2}{Prop_2}, \frac{RRS_1}{Prop_1}) \geq \frac{RRS_i}{Prop_i}$ and thus $u_i(\mathbf{o}) \geq RRS_i$. That is, $LMin_{off}^{Prop}$ with two agents satisfies *RRS*. Finally, due to [11], *RRS* entails $1/2$ -*Prop1*. \square

That is, when there are only two agents, $LMin_{off}^{Prop}$ is (arguably) the “fairest” mechanism among the mechanisms that we study, since it is the only one that satisfies *PO*, $1/2$ -*Prop1*, *MPP*, and *RRS*, simultaneously.

7.1 Borda Valuations

Recall that with Borda valuations, RRS_i is the same for all i . Similarly, $Prop_i$ is the same for all i . Thus, in this setting, $LMin_{off}^{Prop}$

and $LMin_{off}^{RRS}$ are equivalent to $LMin_{off}$. That is, there is effectively a single leximin mechanism, $LMin_{off}$, with Borda valuations.

Clearly, every mechanism that satisfies an efficiency or fairness property without any restriction on the valuation, satisfies the property with Borda valuations. We thus analyze the properties that are not satisfied in the general case, and we begin with *Prop*.

Since $LMin_{off}$ satisfies *MPP* (Theorem 5), then with Borda valuations, $LMin_{off}$ guarantees a 1-additive approximation of *Prop* (Theorem 1). Moreover, we show that $LMin_{off}$ guarantees the best possible constant-factor approximation of *Prop*.

Theorem 7. Even with Borda valuations, there is no mechanism that guarantees an additive constant-factor approximation of *Prop* better than 1.

Proof. Assume by contradiction that there is a mechanism that guarantees an α -additive approximation of *Prop*, with $\alpha < 1$. Since $\alpha < 1$, we can choose n such that $\frac{n-1}{n} > \alpha$. Let $T = 1$ and $m = n$. In addition, assume that the agents use Borda valuations, and let $v_i^1(c_i) = 0$, for any agent i . Given any outcome \mathbf{o} , there must be an agent j that receives a utility of 0 from \mathbf{o} . However, $Prop_j = \frac{n-1}{n}$, and thus the difference between $Prop_j$ and $u_j(\mathbf{o})$ is strictly greater than α . A contradiction. \square

We now show that both RR_{off} and MNW_{off} do not guarantee an additive constant-factor approximation of *Prop*. Specifically, RR_{off} does not guarantee an additive approximation that is better than $m - 1$.

Theorem 8. With Borda valuations, RR_{off} does not guarantee an additive approximation of *Prop* that is better than $m - 1$.

Proof. Assume by contradiction that RR_{off} guarantees an $(m - 1 - \alpha)$ -additive approximation of *Prop*, with $\alpha > 0$. Since $\alpha > 0$, we can choose n such that $\frac{(n-1)(m-1)}{n} > m - 1 - \alpha$. Let $T = n - 1$ and for every round t , and let

$$V^t = \begin{pmatrix} m-1 & m-2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ m-1 & m-2 & \dots & 0 \\ 0 & 1 & \dots & m-1 \end{pmatrix}.$$

Assume that RR_{off} uses the identity permutation as the order over the agents. Clearly, $\mathbf{o} = (c_1, \dots, c_1)$, and thus $u_n(\mathbf{o}) = 0$. On the other hand, $Prop_n = \frac{(n-1)(m-1)}{n}$. That is, the difference between $Prop_n$ and $u_n(\mathbf{o})$ is strictly greater than $m - 1 - \alpha$. A contradiction. \square

Indeed, RR_{off} satisfies *Prop1* [11], and this entails that RR_{off} guarantees an $(m - 1)$ -additive approximation of *Prop*, as we show in Theorem 2. Similarly, since MNW_{off} satisfies *Prop1*, MNW_{off} also guarantees an $(m - 1)$ -additive approximation of *Prop*. On the other hand, MNW_{off} does not guarantee an additive approximation that is better than $(m - 3)/2$.

Theorem 9. Even with Borda valuations, MNW_{off} does not guarantee an additive approximation of *Prop* that is better than $\frac{m-3}{2}$.

Proof sketch. Consider the following example. Let $n = 4, m = 6$, $T = 6$, and for every $t \leq T$,

$$V^t = \begin{pmatrix} 5 & 4 & 3 & 2 & 0 & 1 \\ 0 & 1 & 2 & 3 & 4 & 5 \\ 0 & 1 & 2 & 3 & 4 & 5 \\ 0 & 1 & 2 & 3 & 4 & 5 \end{pmatrix}.$$

MNW_{off} chooses $\mathbf{o} = (c_6, \dots, c_6)$, and $u(\mathbf{o}) = (6, 30, 30, 30)$. In addition, $Prop_i = \frac{30}{4} = 7.5$. Therefore, \mathbf{o} guarantees only a $\frac{3}{2}$ -additive approximation of $Prop$. Our proof generalizes this example, and shows that MNW_{off} does not guarantee an additive approximation of $Prop$ that is better than $\frac{m-3}{2}$. \square

We now analyze $Prop1$. With Borda valuations, $LMin_{off}$ guarantees a 1-additive approximation of $Prop$, and this entails that $LMin_{off}$ satisfies $Prop1$, as we showed in Theorem 2.

We note that Conitzer et al. [11] have raised an open question: is there a mechanism that satisfies PO , $Prop1$, and RRS simultaneously? We partially solve this question- if we restrict the valuations to Borda valuations, then $LMin_{off}$ is such a mechanism, since we showed that it satisfies PO , $Prop1$, and RRS .

Next, we consider MPP , and the negative results for MNW_{off} and RR_{off} still hold. Indeed, Example 1 uses Borda valuations, and it shows that MNW_{off} and RR_{off} do not satisfy MPP . In addition, MNW_{off} still does not satisfy RRS .

Theorem 10. *Even With Borda valuations, MNW_{off} does not satisfy RRS .*

Proof. With Borda valuations, for any agent i , $RRS_i = \sum_{1 \leq t \leq \lfloor T/n \rfloor} \mathbf{cMax}_i(t \cdot n) = \sum_{1 \leq t \leq \lfloor T/n \rfloor} m - 1 = \lfloor \frac{T}{n} \rfloor (m - 1)$. In addition, $Prop_i = \frac{1}{n} \sum_{t=1}^T v_i^t(\mathbf{cMax}_i^t) = \frac{1}{n} \sum_{t=1}^T m - 1 = \frac{T}{n} (m - 1)$. Therefore, if $T \pmod n = 0$, $RRS_i = Prop_i$. Recall the setting in the proof of Theorem 9. The theorem shows that MNW_{off} chooses an outcome, \mathbf{o} , that does not satisfy $Prop$, for any $n > 2$. If $T = n$, $RRS_i = Prop_i$, and thus \mathbf{o} does not satisfy RRS . \square

Finally, we note that RR_{off} still does not satisfy PO .

Theorem 11. *With Borda valuations, RR_{off} does not satisfy PO .*

Overall, $LMin_{off}$ with Borda valuations is the “fairest” mechanism among the mechanisms we study, since it is the only one that satisfies PO , $Prop1$, RRS , MPP , and guarantees a 1-additive approximation of $Prop$.

8 Online Setting

Clearly, all the mechanisms we study do not satisfy the PO efficiency property in the online setting. Moreover, we show that it is impossible to satisfy $Prop1$ or α - RRS for any constant α .

Theorem 12. *There is no online mechanism that satisfies α - RRS , for any constant α .*

Proof. Assume by contradiction that there is an online mechanism that satisfies α - RRS . Let $m = 2, n = 2, T = 3$, and

$$V^1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, X^2 = \begin{pmatrix} \frac{2}{\alpha} & 0 \\ 0 & \frac{\alpha}{2} \end{pmatrix}, Y^2 = \begin{pmatrix} \frac{\alpha}{2} & 0 \\ 0 & \frac{2}{\alpha} \end{pmatrix}$$

$$X^3 = \begin{pmatrix} \frac{4}{\alpha^2} & 0 \\ 0 & \frac{4}{\alpha^2} \end{pmatrix}, Y^3 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

We set V^2 and V^3 according to the decisions of the online mechanism. Specifically, if $o^1 = c_1$ we set $V^2 = X^2$. Otherwise, $V^2 = Y^2$. Next, if $o^1 \neq o^2$ we set $V^3 = X^3$. Otherwise, $V^3 = Y^3$.

Note that for any agent i , $RRS_i \geq \frac{\alpha}{2}$ and thus $\alpha \cdot RRS_i > 0$. Since \mathbf{o} satisfy α - RRS it must be the case that $o^1 \neq o^2$. Otherwise, $V^3 = Y^3$ and thus either $u_1(\mathbf{o}) = 0$ or $u_2(\mathbf{o}) = 0$, which is a

contradiction. Therefore, $o^1 \neq o^2$ and $V^3 = X^3$. In addition, since $0 < \alpha \leq 1$, then $\frac{\alpha}{2} < 1 < \frac{2}{\alpha} < \frac{4}{\alpha^2}$, and thus \mathbf{cMax}_i equals $(1, \frac{2}{\alpha}, \frac{4}{\alpha^2})$ or $(\frac{\alpha}{2}, 1, \frac{4}{\alpha^2})$, for any agent i . That is, $RRS_i \in \{\frac{2}{\alpha}, 1\}$, and thus α - $RRS_i \in \{2, \alpha\}$, for any agent i .

Now, since $o^1 \neq o^2$, there are only 4 possible outcomes:

1. If $\mathbf{o} = (c_1, c_2, c_1)$, $V^2 = X^2$. Thus, $\alpha \cdot RRS_1 = 2$ and $\alpha \cdot RRS_2 = \alpha$. However, $u(\mathbf{o}) = (1 + \frac{4}{\alpha^2}, \frac{\alpha}{2})$. That is, $u_2(\mathbf{o}) < \alpha \cdot RRS_2$.
2. If $\mathbf{o} = (c_2, c_1, c_2)$, $V^2 = Y^2$. Thus, $\alpha \cdot RRS_1 = \alpha$ and $\alpha \cdot RRS_2 = 2$. However, $u(\mathbf{o}) = (\frac{\alpha}{2}, 1 + \frac{4}{\alpha^2})$. That is, $u_1(\mathbf{o}) < \alpha \cdot RRS_1$.
3. If $\mathbf{o} = (c_1, c_2, c_2)$, $V^2 = X^2$. Thus, $\alpha \cdot RRS_1 = 2$ and $\alpha \cdot RRS_2 = \alpha$. However, $u(\mathbf{o}) = (1, \frac{\alpha}{2} + \frac{4}{\alpha^2})$. That is, $u_1(\mathbf{o}) < \alpha \cdot RRS_1$.
4. If $\mathbf{o} = (c_2, c_1, c_1)$, $V^2 = Y^2$. Thus, $\alpha \cdot RRS_1 = \alpha$ and $\alpha \cdot RRS_2 = 2$. However, $u(\mathbf{o}) = (\frac{\alpha}{2} + \frac{4}{\alpha^2}, 1)$. That is, $u_2(\mathbf{o}) < \alpha \cdot RRS_2$.

Therefore, there is no online mechanism that satisfies α - RRS . \square

Next, we show that it is impossible to satisfy $Prop1$. Let $Prop_i^t$ be the $Prop$ value of agent i up to round t , i.e., $Prop_i^t = \sum_{k=1}^{k=t} \frac{v_i^k(\mathbf{cMax}_i^k)}{n}$. Let $vProp1_i^t = \arg \max_{k \in \{1, \dots, t\}} (v_i^k(\mathbf{cMax}_i^k) - v_i^k(o^k))$. It is thus possible to rephrase the definition of $Prop1$: an outcome \mathbf{o} satisfies $Prop1$ if for every agent i , $u_i(\mathbf{o}) + vProp1_i^T \geq Prop_i^T$. Now, let $dProp1_i^t = u_i(o^t) + vProp1_i^t - Prop_i^t$. That is, if $dProp1_i^t \geq 0$ for any i , then \mathbf{o}^t satisfies $Prop1$ up to round t , and if $dProp1_i^t < 0$ for an agent i , then $Prop1$ is not satisfied up to round t . Therefore, intuitively, $dProp1_i^t$ represents the degree in which \mathbf{o}^t satisfies $Prop1$.

Theorem 13. *There is no online mechanism that satisfies $Prop1$.*

Proof sketch. Assume by contradiction that there is an online mechanism that satisfies $Prop1$. We build a scenario with two agents, in which for each even round t , the mechanism must choose an outcome such that $dProp1_i^t$ of one of the agents is the same as her $dProp1_i^{t-2}$, but $dProp1_i^t$ of the other agent is smaller than her $dProp1_i^{t-2}$ by a constant factor. Therefore, after a constant number of rounds, one of the agents has a negative $dProp1_i^t$, which entails that the mechanism does not satisfy $Prop1$. \square

Note that the proofs of Theorems 12 and 13 essentially use the setting of indivisible private goods. Therefore, they strengthen previous results on online fair division [6]. Overall, with no restriction on the valuations, there is no online mechanism that satisfies even the relaxed fairness properties (i.e., RRS or $Prop1$).

8.1 Borda Valuations

We begin with analyzing $Prop$. Our main result here is that $LMin_{on}$ guarantees a 1-additive approximation of $Prop$. To prove this, we first show a unique characteristic of Borda valuations.

Lemma 1. *With Borda valuations, given a round t , and given any vector of n elements, \mathbf{x} , such that (1) every element of \mathbf{x} is a non-negative integer, and (2) the sum of the elements of \mathbf{x} is $m - 1$, there exists a candidate $c \in C$ such that $v_i^t(c) \geq x_i$ for every agent i .*

Proof. Let \mathbf{x} be a vector of n elements, such that (1) every element of \mathbf{x} is a non-negative integer, and (2) the sum of the elements is $m - 1$. For each agent i , let B_i be the set of candidates whose valuations are less than \mathbf{x}_i , i.e., $B_i = \{c \in C \mid v_i^t(c) < \mathbf{x}_i\}$. Let G be a (possibly empty) set of candidates, $G = C \setminus (B_1 \cup \dots \cup B_n)$. Clearly, $|G| \geq m - \sum |B_i|$. Since the agents use Borda valuations, then $|B_i| = \mathbf{x}_i$, and thus $|G| \geq m - \sum \mathbf{x}_i = m - (m - 1) = 1$. That is, a candidate $c \in G$ exists such that $v_i^t(c) \geq \mathbf{x}_i$ for every agent i . \square

Intuitively, Lemma 1 shows that any mechanism can choose an outcome such that the sum of all the agents' valuations is at least $T(m - 1)$. Moreover, a mechanism is able to distribute the valuations arbitrarily, and thus it can choose an outcome that guarantees a utility of $\lfloor \frac{T(m-1)}{n} \rfloor$ for each agent. Therefore, $LMin_{off}$ must also guarantee a utility of $\lfloor \frac{T(m-1)}{n} \rfloor$ for each agent, which means that it satisfies 1-additive approximation of $Prop$. The 1-additive approximation for the online setting is also based on Lemma 1. Intuitively, the lemma shows that for each round, the vector of elements \mathbf{x} can be chosen without knowing the valuations of each agent. Therefore, the vector \mathbf{x} for each round can be chosen greedily, and it thus determines each \mathbf{o}^t in round t such that the utility of \mathbf{o} is at least $\lfloor \frac{T(m-1)}{n} \rfloor$ for each agent.

Next, we show an important property of proportionality in the online setting, i.e., when the candidates of the chosen outcome are determined sequentially. The proof is essentially an extension of Lemma 1 to the online setting. Recall that with Borda valuations, $Prop_i = \frac{T(m-1)}{n}$ for every agent i . Let $Prop^t = \frac{t(m-1)}{n}$, $qProp^t = \lfloor Prop^t \rfloor$, and $rProp^t = (t(m-1)) \bmod n$.

Lemma 2. *With Borda valuations, given a round t , if (1) for every agent i , $u_i(\mathbf{o}^{t-1}) \geq qProp^{t-1}$ and (2) there are at least $rProp^{t-1}$ agents with a utility that is strictly greater than $qProp^{t-1}$, then there exists a candidate c such that (3) for every agent i , $u_i(\mathbf{o}^1, \dots, \mathbf{o}^{t-1}, c) \geq qProp^t$ and (4) there are at least $rProp^t$ agents with a utility that is strictly greater than $qProp^t$.*

We are now ready to prove the approximation result.

Theorem 14. *$LMin_{on}$ guarantees a 1-additive approximation of $Prop$.*

Proof. Let \mathbf{o}^{t-1} be the vector of chosen candidates by $LMin_{on}$ up to round t . If \mathbf{o}^{t-1} satisfies conditions (1) and (2) of Lemma 2, then there exists a candidate c such that conditions (3) and (4) of Lemma 2 hold. By definition, $LMin_{on}$ chooses a candidate \mathbf{o}^t such that $u(\mathbf{o}^t) \succeq u(\mathbf{o}^1, \dots, \mathbf{o}^{t-1}, c)$. That is, \mathbf{o}^t that is chosen by $LMin_{on}$ satisfies conditions (3) and (4). In addition, in the first round $t = 1$ and thus, obviously, $rProp^{t-1} = qProp^{t-1} = 0$. That is, conditions (1) and (2) of Lemma 2 hold when $t = 1$. Therefore, for any round t and agent i , $u_i(\mathbf{o}^t) \geq qProp^t$. Specifically, when $t = T$, $u_i(\mathbf{o}^T) \geq qProp^T = \lfloor Prop_i \rfloor$. By definition, $\lfloor Prop_i \rfloor > Prop_i - 1$, and $\mathbf{o} = \mathbf{o}^T$. That is, $u_i(\mathbf{o}) + 1 > Prop_i$, for any agent i . \square

Similar to the offline setting, the other mechanisms do not guarantee an additive constant-factor approximation of $Prop$. Indeed, RR_{on} is essentially equivalent to RR_{off} when the valuations are restricted to Borda valuations, since for every agent i and in every round t , $cMax_i^t$ are the same (i.e., $m - 1$). That is, RR_{on} does not guarantee an additive approximation that is better than $m - 1$. As for MNW_{on} , since the example in the proof of Theorem 9 holds for any $t \leq T$, MNW_{on} chooses the same outcome as MNW_{off}

in this instance. That is, MNW_{on} does not guarantee an additive approximation of $Prop$ that is better than $\frac{m-3}{2}$.

We now analyze MPP . Unfortunately, MPP cannot be satisfied in the online setting, even with Borda valuations.

Theorem 15. *Even with Borda valuations, there is no online mechanism that satisfies MPP .*

Proof. Assume by contradiction that there is an online mechanism that satisfies MPP . Let $m = 2$, $n = 4$, $T = 2$, and

$$V^1 = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}, X = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}, Y = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

If $V^2 = X$, the only outcome that satisfies $Prop$ is (c_1, c_2) . Therefore, the mechanism must choose the candidate c_1 in the first round. On the other hand, if $V^2 = Y$, the only outcome that satisfies $Prop$ is (c_2, c_1) , and thus the mechanism must choose c_2 in the first round, which is a contradiction. \square

Next, we analyze RRS .

Theorem 16. *With Borda valuations, $LMin_{on}$ satisfies RRS .*

Proof. Let \mathbf{o} be the outcome chosen by $LMin_{on}$. Recall that for any agent i , $RRS_i = \lfloor \frac{T}{n} \rfloor \cdot (m - 1)$ with Borda valuations. According to the proof of Theorem 14, for any agent i , $u_i(\mathbf{o}) \geq qProp^T = \lfloor \frac{T \cdot (m-1)}{n} \rfloor$. Clearly, $\lfloor \frac{T}{n} \rfloor \cdot (m - 1) \geq \lfloor \frac{T}{n} \rfloor \cdot (m - 1)$. We thus get that for any agent i , $u_i(\mathbf{o}) \geq RRS_i$. That is, $LMin_{on}$ satisfies RRS . \square

As for MNW_{on} , since the example in the proof of Theorem 10 holds for $t \leq T$, MNW_{on} chooses the same outcome as MNW_{off} in this instance. That is, MNW_{on} does not satisfy RRS .

Finally, we consider PO . Recall that $LMin_{off}$ and MNW_{off} satisfy PO . In the online setting, even with Borda valuations, they do not satisfy PO anymore.

Theorem 17. *Even with Borda valuations, $LMin_{on}$ and MNW_{on} do not satisfy PO .*

9 Conclusions and Future Work

In this paper, we study collective decision-making mechanisms for a sequence of decisions, when the preferences of the agents are cardinal. However, since we also concentrate on the restriction to Borda valuations, our results are also applicable when the preferences of the agents are ordinal. Overall, we claim that the leximin mechanism, in which the valuations are normalized with the proportional value of each agent, has a significant advantage over the other known mechanisms: in the offline setting, it satisfies PO and MPP , and when there are only two agents it also satisfies $1/2-Prop1$ and RRS . With Borda valuations, the leximin mechanism satisfies $Prop1$ and RRS , and guarantees the best possible constant-factor approximation of $Prop$. Moreover, these results hold both for the offline and online settings.

For future work, we would like to consider other restrictions on the agents' valuations, e.g., valuations that are based on a (given) general scoring rule. We would also like to study what happens when there are many rounds. Specifically, we conjecture that given any n and m , there is a t such that the leximin mechanism with Borda valuations satisfies $Prop$ for any $T \geq t$. Finally, since it is impossible to achieve proportionality or its relaxations in the online setting, we would like to extend our framework and analyze probabilistic mechanisms.

Acknowledgements

This research has been partly supported by the Ministry of Science, Technology & Space (MOST), Israel.

References

- [1] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh, 'On-line fair division: Analysing a food bank problem', in *Proceedings of the Twenty-Fourth International Conference on Artificial Intelligence (IJCAI-2015)*, pp. 2540—2546, (2015).
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris, 'Fair division of indivisible goods: A survey', in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pp. 5385–5393, (2022).
- [3] Haris Aziz and Nisarg Shah, 'Participatory budgeting: Models and approaches', in *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, 215–236, Springer, (2021).
- [4] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah, 'Proportionally fair online allocation of public goods with predictions', *arXiv preprint arXiv:2209.15305*, (2022).
- [5] Dorothea Baumeister, Sylvain Bouveret, Jérôme Lang, Nhan-Tam Nguyen, Trung Thanh Nguyen, and Jörg Rothe, 'Scoring rules for the allocation of indivisible goods', in *Proceedings of the Twenty-First European Conference on Artificial Intelligence (ECAI-2014)*, pp. 75–80, (2014).
- [6] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas, 'How to make envy vanish over time', in *Proceedings of the 2018 ACM Conference on Economics and Computation (EC-2018)*, pp. 593–610, (2018).
- [7] Sylvain Bouveret and Michel Lemaître, 'Characterizing conflicts in fair division of indivisible goods using a scale of criteria', *Autonomous Agents and Multi-Agent Systems*, **30**(2), 259–290, (2016).
- [8] Robert Bredereck, Till Fluschnik, and Andrzej Kaczmarczyk, 'When votes change and committees should (not)', in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pp. 144–150, (2022).
- [9] Markus Brill, Paul Gözl, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker, 'Approval-based apportionment', *Mathematical Programming*, 1–29, (2022).
- [10] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon, 'Justified representation for perpetual voting', *IEEE Access*, **9**, 96598–96612, (2021).
- [11] Vincent Conitzer, Rupert Freeman, and Nisarg Shah, 'Fair public decision making', in *Proceedings of the 2017 ACM Conference on Economics and Computation (EC-2017)*, pp. 629–646, (2017).
- [12] Andreas Darmann and Christian Klamler, 'Proportional borda allocations', *Social Choice and Welfare*, **47**(3), 543–558, (2016).
- [13] Andreas Darmann and Joachim Schauer, 'Maximizing nash product social welfare in allocating indivisible goods', *European Journal of Operational Research*, **247**(2), 548–559, (2015).
- [14] Brandon Fain, Kamesh Munagala, and Nisarg Shah, 'Fair allocation of indivisible public goods', in *Proceedings of the 2018 ACM Conference on Economics and Computation (EC-2018)*, pp. 575–592, (2018).
- [15] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon, 'Multiwinner voting: A new challenge for social choice theory', in *Trends in computational social choice*, chapter 2, 27–47, AI Access, (2017).
- [16] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer, 'Fair social choice in dynamic settings', in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pp. 4580–4587, (2017).
- [17] Jugal Garg, Pooja Kulkarni, and Aniket Murhekar, 'On fair and efficient allocations of indivisible public goods', in *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS-2021)*, pp. 22:1–22:19, (2021).
- [18] Jiafan He, Ariel D. Procaccia, CA Psomas, and David Zeng, 'Achieving a fairer future by changing the past', in *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*, pp. 343–349, (2019).
- [19] Ido Kahana and Noam Hazon, 'The leximin approach for a sequence of collective decisions', *arXiv preprint arXiv:2305.18024*, (2023).
- [20] Ian Kash, Ariel D. Procaccia, and Nisarg Shah, 'No agent left behind: Dynamic fair division of multiple resources', *Journal of Artificial Intelligence Research*, **51**, 579–603, (2014).
- [21] Martin Lackner, 'Perpetual voting: Fairness in long-term decision making', in *Proceedings of the AAAI conference on artificial intelligence (AAAI-2020)*, pp. 2103–2110, (2020).
- [22] Martin Lackner and Jan Maly, 'Proportional decisions in perpetual voting', in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-2023)*, pp. 5722–5729, (2023).
- [23] Martin Lackner, Jan Maly, and Simon Rey, 'Fairness in long-term participatory budgeting', in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2020)*, pp. 1566–1568, (2021).
- [24] Martin Lackner and Piotr Skowron, *Multi-winner voting with approval preferences*, Springer Nature, 2023.
- [25] Jérôme Lang and Lirong Xia, 'Voting in combinatorial domains', in *Handbook of computational social choice*, eds., Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, chapter 9, 197–222, Cambridge University Press, (2016).
- [26] Tyler Lu and Craig Boutilier, 'Budgeted social choice: From consensus to personalized decision making', in *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence (IJCAI-2021)*, pp. 280–286, (2011).
- [27] Simon Rey and Jan Maly, 'The (computational) social choice take on indivisible participatory budgeting', *arXiv preprint arXiv:2303.00621*, (2023).
- [28] Piotr Skowron and Adrian Górecki, 'Proportional public decisions', in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-2022)*, pp. 5191–5198, (2022).