Free-Riding in Multi-Issue Decisions

Martin Lackner
DBAI, TU Wien
Vienna, Austria
lackner@dbai.tuwien.ac.at

Jan Maly
ILLC, University of Amsterdam
Amsterdam, The Netherlands
j.f.maly@uva.nl

Oliviero Nardi
DBAI, TU Wien
Vienna, Austria
oliviero.nardi@tuwien.ac.at

ABSTRACT

Voting in multi-issue domains allows for compromise outcomes that satisfy all voters to some extent. Such fairness considerations, however, open the possibility of a special form of manipulation: free-riding. By untruthfully opposing a popular opinion in one issue, voters can receive increased consideration in other issues. We study under which conditions this is possible. Additionally, we study free-riding from a computational and experimental point of view. Our results show that free-riding in multi-issue domains is largely unavoidable, but comes at a non-negligible individual risk for voters. Thus, the allure of free-riding is smaller than one could intuitively assume.

KEYWORDS

voting; strategic aspects; multi-issue elections; free-riding

ACM Reference Format:
Martin Lackner, Jan Maly, and Oliviero Nardi. 2023. Free-Riding in Multi-Issue Decisions. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Elections are a fundamental and well-studied form of collective decision making. One can often observe that elections do not occur as isolated events with a tightly constrained decision space (i.e., only a small number of candidates). Instead, a group of voters needs to make several decisions, either at the same time (cf. multiple referenda [3, 6, 8] or voting over combinatorial domains [33]) or over time (cf. perpetual voting [14, 26] or successive committees [11]). For example, the council of a faculty or the members of a sports club have to make several independent decisions each year. By considering these individual decisions in conjunction, one can achieve more equitable outcomes than would otherwise be possible. As the combinatorial complexity increases with the number of issues, so does the possibility of finding good comprise outcomes.

However, by striving for fairness across multiple issues, we open the door to a specific, simple form of manipulation: *free-riding*. We define free-riding as untruthfully opposing a necessarily winning candidate. That is, if there is a very popular (maybe unanimous) candidate for a certain issue, it typically does not change the outcome if one voter does not approve this candidate. Under the assumption that the voting rule in use tries to establish some form of fairness, it will give this voter additional consideration as she does not approve the choice in this issue. As we show in our paper, this form of manipulation is possible almost universally in multi-issue voting.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

The problem of free-riding is particularly apparent if issues are decided sequentially. Then, presented with a popular candidate that is certain to win, a voter may be especially tempted to misrepresent her preferences. This is because untruthfully opposing a winning candidate artificially lowers the voter's (calculated) satisfaction and thus gives the voter additional weight for subsequent issues, if the voting rule is taking past satisfaction into account. Thus, intuitively, it may appear as if free-riding is a form of risk-free manipulation.

The main contribution of our paper is to refute this intuition. While free-riding is indeed often a successful form of manipulation, it is far from trivially beneficial for free-riding individuals. For our analysis, we consider two fundamentally different categories of voting rules: rules based on a global optimization problem and rules based on sequential decisions. Within both categories, we consider voting rules based on order-weighted averages (OWA [3, 45]) and on Thiele scores (inspired by multiwinner voting [31, 44]). Based on these classes, we obtain the following results:

- First, we show that almost every OWA and Thiele rule as well as their sequential counterparts are susceptible to free-riding. The utilitarian rule, maximizing the sum of utilities, is the only exception.
- Unsurprisingly, it is computationally hard to determine the outcome for OWA and Thiele rules based on global optimization. However, we show even stronger hardness results: even when the winner of an issue is known, it remains computationally hard to determine whether free-riding in this issue is possible. Thus, for rules based on global optimization, it is computationally difficult and may require full information to free-ride. We conclude from these results that for optimization-based rules, free-riding is at least no more of a concern than the general problem of strategic voting.
- For sequential OWA and Thiele rules, we observe an interesting phenomenon. Here, it may be that free-riding in an issue leads to a *lower* satisfaction in subsequent issues. Thus, free-riding for these voting rules is not risk-free. Moreover, we show that it is a computationally hard task to determine whether free-riding is beneficial. We note that this decision requires full preference information about all issues; in the case of incomplete information voters cannot determine the impact of free-riding.
- Finally, voters might still decide to free-ride without certainty about the outcome if the risk is small enough. To study this question, we complement our theoretical analysis with numerical simulations to quantify this risk. Our simulations show that the risk of free-riding is indeed significant, even though positive outcomes are more likely.

In general, our results show that free-riding in multi-issue voting is not as simple and risk-free as one could intuitively assume.

1.1 Related Work

Our work falls in the broad class of *voting in combinatorial domains* [33]. In contrast to many works in this field (e.g., [1, 8, 9, 16, 32]), we assume that voters' preferences are separable (i.e., independent) between issues.

Our work is most closely related to papers on multiple referenda. Amanatidis et al. [3] study the computational complexity of OWA voting rules in multiple referenda, including questions of strategic voting. In a similar model, Barrot et al. [6] consider questions of manipulability: how does the OWA vector impact the susceptibility to manipulation. In contrast to our paper, these two papers do not consider free-riding. We discuss more technical connections between these papers and ours later in the text.

Another related formalism is perpetual voting [26], which essentially corresponds to voting on multi-issue decisions in sequential order. In this setting, issues are chronologically ordered, i.e., decided one after the other. The work of Lackner [26] and its follow-up by Lackner and Maly [27, 28] do not consider strategic issues. Further, Bulteau et al. [14] move to a non-sequential (offline) model of perpetual voting and study proportional representation in this setting.

A third related formalism is that of *public decision making* [15]. As in our model, public decision making considers k issues and for each one alternative has to be chosen. This model is more general than ours in that it allows arbitrary additive utilities (whereas we consider only binary utilities, i.e., approval ballots). Our works differ in that Conitzer et al. [15] focus on fairness properties, whereas our focus is on strategic aspects. Fairness considerations in public decision making have further been explored by Skowron and Górecki [43]. Note that both papers [15, 43] assume that all issues are decided in parallel (offline) – in contrast to perpetual voting [26].

Our model is also related to multi-winner voting [20, 31]. The main difference is that instead of selecting k candidates from the same set of candidates, we have individual candidates for each of the k issues. In our paper, we adapt the class of Thiele rules from the multi-winner setting to ours. This class has been studied extensively, both axiomatically [4, 30, 39, 40] and computationally [5, 13, 23, 42]. The concept of free-riding has also been considered for multi-winner elections [7, 36, 41]. Here, free-riding refers to "subset-manipulation", i.e., to submit only a subset of one's truly approved candidates. We note that this notion of free-riding is related to ours in its essence, but technically distinct.

In multi-winner voting, there is also substantial literature on the relationship between fairness (often proportionality) and strategyproofness, e.g., [17, 25, 29, 34, 36].

Finally, free-riding is a very general phenomenon and has been widely studied in the economic literature on public goods [24, 38]. It has also been considered in more technical domains, such as free-riding in memory sharing [21].

2 THE MODEL

As is customary, we write $[k]$ to denote $\{1, \dots, k\}$.

We study a form of multi-issue decision making, where for each issue there are two or more possible options available. Furthermore, we assume that for each issue each voter submits an approval ballot, i.e., a subset of candidates that she likes. Formally, k denotes the number of issues and $C_1, \dots, C_k$ the respective sets of candidates. Let $N = [n]$ denote the set of voters. We write $A_i(v) \subseteq C_i$ for the approval ballot of voter v concerning issue i. In combination, we call such a triple $\mathcal{E} = (\{C_i\}_{i\in[k]}, N, \{A_i\}_{i\in[k]})$ an *election*. If k is clear from the context, we write $\bar{C}$ for $\{C_i\}_{i\in[k]}$ and $\bar{A}$ for $\{A_i\}_{i\in[k]}$.

An *outcome* of an election is a k-tuple $\bar{w} = (w_1, \dots, w_k)$ with $w_i \in C_i$. Given an election $\mathcal{E}$ and an outcome $\bar{w}$, the *satisfaction* of voter $v \in N$ with $\bar{w}$ is $sat_{\mathcal{E}}(v, \bar{w}) = |\{1 \le i \le k : w_i \in A_i(v)\}|$ In other words, the satisfaction of a voter is the number of issues that were decided in this voter's favour.[1] Furthermore, we write $s_{\mathcal{E}}(\bar{w}) = (s_1, \dots, s_n)$ to denote the n-tuple of satisfaction scores $(sat_{\mathcal{E}}(v, \bar{w}))_{v\in N}$ sorted in increasing order, i.e., $s_1 \le s_2 \le \dots \le s_n$. If the election $\mathcal{E}$ is clear from the context, we omit it in the notation.

There are two main voting rules that have been studied in this setting: maximizing the total satisfaction and maximizing the satisfaction of the least satisfied voter.[2]

- The *utilitarian rule* returns an outcome $\bar{w}$ that maximizes $\sum_{v\in N} sat(v, \bar{w})$. This rule corresponds to selecting issue-wise the candidate with the most approvals.
- The *egalitarian rule* returns an outcome $\bar{w}$ that maximizes $\min_{v\in N} sat(v, \bar{w})$.

The egalitarian rule is NP-hard to compute [3], while the utilitarian rule is computable in polynomial time (as one can decide each issue separately).

The egalitarian rule has the disadvantage that often many outcomes are optimal in the egalitarian sense. In such cases, it would be desirable to also pay attention to the second-least satisfied voter, third-least, etc. This leads to the *leximin* rule.

- The *leximin rule* is based on the leximin ordering $\succ$. Given two outcomes $\bar{w}$ and $\bar{w}'$ with $s(\bar{w}) = (s_1, \dots, s_n)$ and $s(\bar{w}') = (s'_1, \dots, s'_n)$, $\bar{w} \succ \bar{w}'$ if there exists an index $j \in [n]$ such that $s_1 = s'_1, \dots, s_{j-1} = s'_{j-1}$ and $s_j > s'_j$. The *leximin rule* returns an outcome $\bar{w}$ that is maximal with respect to $\succ$.

Example 1. *Consider an election with 100 voters and 4 issues with the same three candidates, $\{a, b, c\}$. There are 66 voters that approve $\{a\}$ in all issues, 33 voters that approve $\{b\}$ in all issues, and one voter approves always $\{c\}$. The utilitarian rule selects the outcome $\bar{w}_1 = (a, a, a, a)$ as it achieves a total satisfaction of $\sum_{v\in N} sat(v, \bar{w}_1) = 4 \cdot 66$. The leximin rule selects $\bar{w}_2 = (a, a, b, c)$ (or a permutation thereof) with $s(\bar{w}_2) = (\underbrace{1, \dots, 1}_{34\ times}, \underbrace{2, \dots, 2}_{66\ times})$. The egalitarian rule can select any outcome that contains a, b, and c at least once, including the rather questionable outcome $\bar{w}_3 = (a, b, c, c)$ with $s(\bar{w}_3) = (1, \dots, 1, 2)$.*

[1] Note that different notions of satisfaction are possible; for instance, we could assume that the voters have fine-grained preferences over the issues and the candidates. However, our simple model is a natural starting point, and we leave the investigation of different notions of satisfaction as future work.

[2] These two voting rules (in the context of binary elections) are referred to as minsum and minimax by Amanatidis et al. [3]. Note that in the case of binary elections, the satisfaction of a voter v with $\bar{w}$ corresponds to k minus the Hamming distance (symmetric difference) between $\{i \in [k] : A_i(v) = \{1\}\}$ and $\{i \in [k] : w_i = 1\}$. The minsum rule minimizes the sum of Hamming distances; the minimax rule minimizes the maximum Hamming distance. This is equivalent to our approach of maximizing the total or minimum satisfaction.

2.1 Optimization-Based Rules

In the following, we describe two classes of multi-issue voting rules based on maximizing scores. OWA voting rules for multi-issue domains were proposed by Amanatidis et al. [3] and are based on *ordered weighted averaging* operators [45]. An OWA voting rule is defined by a set of vectors $\{\alpha^n\}_{n\geq 1}$, where each $\alpha^n = (\alpha_1, \dots, \alpha_n)$ has length n and satisfies $\alpha_1 > 0$ and $\alpha_j \geq 0$ for $j \in [n]$. Given an election with n voters, the score of an outcome $\bar{w}$ subject to α^n is

$$OWA_{\alpha^n}(\bar{w}) = \alpha^n \cdot s(\bar{w}),$$

where $\cdot$ is the scalar (dot) product. The OWA rule returns an outcome with maximum OWA_{α^n}-score. If more than one outcome achieves the maximum score, we use a fixed tie-breaking order among outcomes. We typically omit the superscript of α^n, as n is clear from the context.

Note that the utilitarian rule corresponds to $\alpha^n = (1/n, \dots, 1/n)$, the egalitarian rule corresponds to $\alpha^n = (1, 0, \dots, 0)$, and the leximin rule to $\alpha^n = (1, 1/kn, 1/k^2n^2, \dots)$.[3]

Proposition 1. *The OWA rule defined by* $\alpha = (1, \frac{1}{kn}, \frac{1}{k^2n^2}, \dots)$ *is equivalent to the leximin rule.*

The second class is based on Thiele methods (introduced by Thiele [44], see the survey by Lackner and Skowron [31]). While Thiele methods are a class of multi-winner voting rules, they can be adapted to our setting straightforwardly. A voting rule in the Thiele class is defined by a function $f\colon \mathbb{N} \to \mathbb{R}^{\geq 0}$ satisfying $f(1) > 0$ and $f(i) \geq f(i+1)$ for all $i \in \mathbb{N}$. The f-Thiele rule assigns a score of

$$\mathit{Thiele}_f(\bar{w}) = \sum_{v \in N} \sum_{i=1}^{sat(v,\bar{w})} f(i)$$

to an outcome $\bar{w}$ and returns an outcome with maximum score. Intuitively, these are weighted approval rules for which the weight assigned to each voter only depends on her satisfaction. Note that the utilitarian rule corresponds to $f_{util}(i) = 1$. The egalitarian and leximin rules do not appear in this class.[4] Another important Thiele rule is $f(i) = 1/i$, which is called *Proportional Approval Voting* in the multi-winner setting. We also refer to this Thiele rule as PAV.

Example 2. *Continuing with the election of Example 1, we see that PAV selects* $\bar{w}_4 = (a, a, a, b)$ *(or a permutation thereof) with* $\mathit{Thiele}_{PAV}(\bar{w}_4) = 66 + 33 + 22 + 33$. *Note that PAV is more majoritarian than leximin as it essentially ignores the single* $\{c\}$*-voter.*

2.2 Sequential Rules

We also consider sequential variants of both the OWA and Thiele classes. Sequential rules construct the outcome in *rounds*, one issue after the other. To define them, we require an ordering over issues. In this paper, we make no assumptions about the origin of these orderings, but a natural order may follow from time (issues are decided at different points in time)[5] or importance (important decisions are made first). The advantage of sequential rules is that they are computable in polynomial time. They can be viewed as approximation algorithms of their optimization-based counterparts.

[3]This definition requires the assumption of a fixed number of issues k.

[4]However, if we fix n and k, leximin can be "simulated" by, e.g., $f_{lex}(i) = 1/(kn)^{i-1}$.

[5]This corresponds to the model of perpetual voting [26], where a sequence of collective decisions has to be made at different points in time.

To formally define sequential rules, we assume that issues are decided in order $1, \dots, k$. The sequential α-OWA rule is defined as follows: If $w_1, \dots, w_{i-1}$ are already selected for issues $1, \dots, i-1$, then we select for issue i a candidate $c \in C_i$ that maximizes $OWA_\alpha(w_1, \dots, w_{i-1}, c)$. This is repeated until all issues have been decided. Similarly, for sequential f-Thiele we iteratively choose for issue i a candidate $c \in C_i$ that maximizes $\mathit{Thiele}_f(w_1, \dots, w_{i-1}, c)$.

2.3 Free-Riding

In this paper, we study a specific form of strategic manipulation called *free-riding*. Intuitively, this means that a voter misrepresents her preferences on an issue where her favourite candidate wins also without her support. If the used voting rule takes the satisfaction of voters into account (as most OWA and Thiele methods do), such a manipulation can increase the voter's influence on other issues.

Example 3. *Consider an election with three voters and two issues. The first issue is uncontroversial: all voters approve candidate a. The second issue is highly controversial: all voters approve different candidates* ($A_2(1) = \{x\}$, $A_2(2) = \{y\}$, $A_2(3) = \{z\}$). *If the egalitarian rule (with some tie-breaking) is used to determine the outcome, it could select, e.g., the outcome* (a, x). *This leaves voters* 2 *and* 3 *less satisfied than voter* 1. *Both of them could free-ride to improve their satisfaction. Consider voter* 2. *If voter* 2 *changes her ballot on the first issue to another candidate, the outcome changes to* (a, y) *as it gives all voters a satisfaction of* 1 *(according to their ballots). As voter* 2*'s true preferences are positive towards a, this manipulation was successful.*

In the following, given an election $\mathcal{E}$ and a rule $\mathcal{R}$ such that $\mathcal{R}(\mathcal{E}) = (w_1, \dots, w_k)$, we indicate w_i as $\mathcal{R}(\mathcal{E})_i$.

Definition 1. *Consider an election* $\mathcal{E} = (\{C_i\}_{i\in[k]}, N, \{A_i\}_{i\in[k]})$, *a voter* $v \in N$ *and a voting rule* $\mathcal{R}$. *Let* $\mathcal{R}(\mathcal{E}) = (w_1, \dots, w_k)$. *We say that voter* v *can free-ride in election* $\mathcal{E}$ *on issues* $I \subseteq [k]$ *if there exists another election* $\mathcal{E}^* = (\{C_i\}_{i\in[k]}, N, \{A_i^*\}_{i\in[k]})$ *that only differs from* $\mathcal{E}$ *in the approvals of* v *for issues in* I *such that, for all* $i \in I$, $w_i \in A_i(v)$, $w_i \notin A_i^*(v)$ *and* $\mathcal{R}(\mathcal{E}^*)_i = w_i$. *In this case, we also say that* v *can free-ride in* $\mathcal{E}$ *via* $\mathcal{E}^*$.

Usually, we say a voter can manipulate if she can achieve a higher satisfaction by misrepresenting her preferences. In contrast, Definition 1 makes no assumptions about the satisfaction of the free-riding voter. Instead, we only require that the manipulator can misrepresent her preference in an issue without changing the outcome of the issue. This might lead to the same, a higher or lower satisfaction for the manipulator. This distinction will be crucial when talking about the risk of free-riding.

We will also sometimes consider a more general notion of free-riding. Here, we lift the constraint that the outcome on the issues where free-riding occurs remains exactly the same. We just require that the new winning candidate is still (truthfully) approved by the manipulator. At its core, generalized free-riding is based on the assumption that voters are indifferent between approved candidates. To define *generalized free-riding* formally, we replace $\mathcal{R}(\mathcal{E}^*)_i = w_i$ in Definition 1 with $\mathcal{R}(\mathcal{E}^*)_i \in A_i(v)$.

Finally, we say that *a voting rule* $\mathcal{R}$ *can be manipulated by (generalized) free-riding* if there exists an election $\mathcal{E}$, a voter v and an election $\mathcal{E}^*$ such that v can perform (generalized) free-riding in $\mathcal{E}$ via $\mathcal{E}^*$ and $sat_{\mathcal{E}}(v, \mathcal{R}(\mathcal{E})) < sat_{\mathcal{E}}(v, \mathcal{R}(\mathcal{E}^*))$.

3 POSSIBILITY AND RISK OF FREE-RIDING

In this section, we study for which voting rules free-riding is possible and under which conditions it is safe (in the sense that it cannot lead to a decrease in the satisfaction of the free-riding voter). Firstly, we observe that the results for different issues do not influence each other for the utilitarian rule, hence free-riding on one issue has no effect on the outcome of other issues. Therefore, the utilitarian rule cannot be manipulated by (generalized) free-riding.

Proposition 2. *The utilitarian rule cannot be manipulated by (generalized) free-riding.*

However, it turns out that every other rule in the classes we study can be manipulated by free-riding.

Theorem 3. *Every (sequential) Thiele and (sequential) OWA rule except the utilitarian rule can be manipulated by free-riding.*

Proof. Let $\mathcal{R}$ be an OWA-Rule that is not the utilitarian rule. Then there exists a k for which the vector α for k voters satisfies $\alpha_1 > \alpha_k$. Clearly, $k \geq 2$. Consider an election with 2 issues and k voters. In each issue there are k candidates $a_1, \dots a_k$. In the first issue, voters 1 and 2 approve a_1. Every other voter $i \in \{3, \dots, k\}$ approves a_i. In the second issue voter 1 approves a_1, voter 2 approves a_2 and all other voters approve both a_1 and a_2. We assume that candidates with a lower index are preferred by the tie-breaking, which is applied lexicographically. Selecting a candidate other than a_1 in the first issue leads to satisfaction vector $(0, 1, \dots, 1, 2)$, independently of whether a_1 or a_2 is selected in issue 2. On the other hand, selecting a_1 in issue 1 leads to satisfaction vector $(1, 1, \dots, 1, 2)$, independently of whether a_1 or a_2 is selected in issue 2. This means (a_1, a_1) and (a_1, a_2) lead to the highest OWA score. By tie-breaking, (a_1, a_1) wins. Now, we claim that voter 2 can free-ride by approving a_2 instead of a_1 in the first issue. Assume first, that a candidate other than a_1 or a_2 is selected in the first issue. This still leads to the same satisfaction vector independently of whether a_1 or a_2 is selected in issue 2. However choosing a_1 in both issues now leads to the vector $(0, 1, \dots, 1, 2)$. Choosing a_1 in issue 1 and a_2 in issue 2 leads satisfaction 1 for every voter. Choosing a_2 both times or first a_2 and then a_1 is symmetric. As $\alpha_1 > \alpha_k$ we know that

$$\alpha \cdot (1, \dots, 1) = \sum_{i=1}^{k} \alpha_i > \alpha_k - \alpha_1 + \sum_{i=1}^{k} \alpha_i = \alpha \cdot (0, 1, \dots, 1, 2)$$

It follows that (a_1, a_2) and (a_2, a_1) are the outcomes maximizing the OWA score. By tie-breaking, (a_1, a_2) is the winning outcome. It follows that voter 2 did successfully free-ride.

The proofs for sequential OWA and (sequential) Thiele rules are similar. Details can be found in the full version of the paper. □

Hence, free-riding is essentially unavoidable if we want to guarantee fairer outcomes using Thiele or OWA rules. Intuitively, free-riding seems to offer a simple and risk-free way to manipulate. And indeed, it is risk-free for some voting rules, such as the leximin rule.

Proposition 4. *Free-riding cannot reduce the satisfaction of the free-riding voter when the leximin rule is used, but it can increase the satisfaction of the free-riding voter.*

Proof. It follows directly from Theorem 3 that free-riding can increase the satisfaction of the free-riding voter. Let us show that it can never decrease the satisfaction of the free-riding voter. Let $\mathcal{E}$ be an election, $\bar{w}$ be the outcome of $\mathcal{E}$ under the leximin rule, and consider a voter v^* such that v^* can free-ride in issue k. Finally, let $\mathcal{E}^*$ be the election after v^* free-rides and $\bar{w}^*$ the outcome of $\mathcal{E}^*$ under the leximin rule. In the following we write $N_i^{\mathcal{E}}(\bar{w}) = \{v \in N \mid sat_{\mathcal{E}}(v, \bar{w}) = i\}$. Now, as v^* free-rides, i.e., the winner in issue k is the same in $\bar{w}$ and $\bar{w}^*$, we know that v^* approves the winner of issue k in her honest ballot in $\mathcal{E}$ and does not approve the winner of issue k in her free-riding ballot in $\mathcal{E}^*$. It follows the satisfaction of v^* with $\bar{w}^*$ resp. $\bar{w}$ in $\mathcal{E}$ is higher by one than in $\mathcal{E}^*$, i.e.,

$$\begin{aligned} sat_{\mathcal{E}^*}(v^*, \bar{w}^*) &= sat_{\mathcal{E}}(v^*, \bar{w}^*) - 1 \quad \text{as well as} \\ sat_{\mathcal{E}^*}(v^*, \bar{w}) &= sat_{\mathcal{E}}(v^*, \bar{w}) - 1. \end{aligned} \tag{1}$$

All other voters submit the same ballot in $\mathcal{E}$ and $\mathcal{E}^*$. Hence, for all $v \neq v^*$ we have

$$\begin{aligned} sat_{\mathcal{E}^*}(v, \bar{w}^*) &= sat_{\mathcal{E}}(v, \bar{w}^*) \quad \text{as well as} \\ sat_{\mathcal{E}^*}(v, \bar{w}) &= sat_{\mathcal{E}}(v, \bar{w}). \end{aligned} \tag{2}$$

Now assume for the sake of a contradiction that $sat_{\mathcal{E}}(v^*, \bar{w}) > sat_{\mathcal{E}}(v^*, \bar{w}^*)$, i.e., free-riding led to a lower satisfaction for v^* with respect to her honest ballot.

As $\bar{w}^*$ is the winning outcome of $\mathcal{E}^*$, we know that $\bar{w}^* > \bar{w}$ according to the leximin order in $\mathcal{E}^*$. In other words, there is a j such that $|N_j^{\mathcal{E}^*}(\bar{w}^*)| < |N_j^{\mathcal{E}^*}(\bar{w})|$ and $|N_\ell^{\mathcal{E}^*}(\bar{w}^*)| = |N_\ell^{\mathcal{E}^*}(\bar{w})|$ for all $\ell < j$. We claim that the deciding index j cannot be smaller than $sat_{\mathcal{E}^*}(v^*, \bar{w}^*)$ as for all smaller indices $\ell < sat_{\mathcal{E}^*}(v^*, \bar{w}^*)$ it follows from $sat_{\mathcal{E}}(v^*, \bar{w}) > sat_{\mathcal{E}}(v^*, \bar{w}^*)$ that v^* is not in $N_\ell^{\mathcal{E}}(\bar{w}^*)$, $N_\ell^{\mathcal{E}^*}(\bar{w}^*)$, $N_\ell^{\mathcal{E}}(\bar{w})$ and $N_\ell^{\mathcal{E}^*}(\bar{w})$. Therefore, it follows from (2) that $|N_\ell^{\mathcal{E}}(\bar{w}^*)| = |N_\ell^{\mathcal{E}^*}(\bar{w}^*)|$ and $|N_\ell^{\mathcal{E}}(\bar{w})| = |N_\ell^{\mathcal{E}^*}(\bar{w})|$. Hence, $j < sat_{\mathcal{E}^*}(v^*, \bar{w}^*)$ would be a contradiction to the assumption that $\bar{w}$ is the leximin outcome of $\mathcal{E}$ and hence leximin preferred to $\bar{w}^*$ in $\mathcal{E}$.

Therefore, we know that $|N_\ell^{\mathcal{E}^*}(\bar{w}^*)| = |N_\ell^{\mathcal{E}^*}(\bar{w})|$ for all $\ell < sat_{\mathcal{E}^*}(v^*, \bar{w}^*) \leq j$ and

$$|N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w}^*)| \leq |N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w})|.$$

It follows that also $|N_\ell^{\mathcal{E}}(\bar{w}^*)| = |N_\ell^{\mathcal{E}^*}(\bar{w}^*)| = |N_\ell^{\mathcal{E}^*}(\bar{w})| = |N_\ell^{\mathcal{E}}(\bar{w})|$ for all $\ell < sat_{\mathcal{E}^*}(v^*, \bar{w}^*) \leq j$. Finally, it follows from (1) that v^* is in $N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w}^*)$ but not in $N^{\mathcal{E}}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w}^*)$. Moreover, because we assumed $sat_{\mathcal{E}}(v^*, \bar{w}) > sat_{\mathcal{E}}(v^*, \bar{w}^*)$, v^* is neither in $N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w})$ nor in $N^{\mathcal{E}}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w})$. Therefore, we have

$$\begin{aligned} |N^{\mathcal{E}}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w}^*)| + 1 = |N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w}^*)| \leq \\ |N^{\mathcal{E}^*}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w})| = |N^{\mathcal{E}}_{sat_{\mathcal{E}^*}(v^*, \bar{w}^*)}(\bar{w})|. \end{aligned}$$

However, that means that $\bar{w}^*$ is leximin preferred to $\bar{w}$ in $\mathcal{E}$, which is a contradiction to the assumption that $\bar{w}$ is the outcome of $\mathcal{E}$. □

It remains an open problem to generalize this result to other rules based on global optimization. However, we observe that for most sequential voting rules, free-riding may lead to a decrease in satisfaction. First, we can show that this holds for all sequential Thiele rules, except the utilitarian rule.

Proposition 5. *Let $f : \mathbb{N} \to \mathbb{R}^{\geq 0}$ be a function for which there is an $i \in \mathbb{N}$ such that $f(i) > f(i+1)$. Then, under the sequential f-Thiele rule, free-riding can reduce the satisfaction of the free-riding voter.*

Proof. Consider a sequential f-Thiele rule such that $f(i) > f(i+1)$ and consider the following election with nine voters, $i+4$ issues and candidates $a, \ldots, g$ for all issues. We assume alphabetic tie-breaking. The approvals for each issue are given by these tuples:

$$A_1 = \cdots = A_{i-1} = (\{a\}, \{a\}, \{a\}, \{a\}, \{a\}, \{a\}, \{a\}, \{a\}, \{a\})$$
$$A_i = (\{a\}, \{a\}, \{a\}, \{b\}, \{b\}, \{c\}, \{d\}, \{e\}, \{f\})$$
$$A_{i+1} = (\{b\}, \{a\}, \{c\}, \{b\}, \{b\}, \{a\}, \{a\}, \{a\}, \{b\})$$
$$A_{i+2} = (\{b\}, \{a\}, \{c\}, \{b\}, \{e\}, \{a\}, \{f\}, \{a, b\}, \{g\})$$
$$A_{i+3} = (\{b\}, \{c\}, \{d\}, \{e\}, \{b\}, \{f\}, \{a\}, \{a\}, \{g\})$$

Then, $\{a\}$ is clearly the winner for the first $i-1$ issues. Thus, all voters have a satisfaction of $i-1$ before issue i and a wins on issue i as it has the most supporters. In issue $i+1$, a and b increase the Thiele score by $f(i+1) + 3f(i)$, while c increases the score only by $f(i+1)$. By tie-breaking, a wins again. Then, in issue $i+2$, a increases the score by $f(i+2)+2f(i+1)$, b by $f(i)+2f(i+1)$, while all other candidates by at most $f(i)$. As $f(i) > f(i+1) \geq f(i+2)$ (together with the tie-breaking rule if $f(i+1) = 0$), it follows that b wins in issue $i+2$. Finally, in issue $i+3$, a increases the score by $f(i+1) + f(i+2)$, b by $f(i) + f(i+2)$ and all other candidates increase the score by at most $f(i)$. Hence b wins.

Now assume voter 1 changes her preferences and free-rides in issue i. It is straightforward to check that a remains the winner for issue i, but winner in issue $i+1$ changes to b while a now wins for issue $i+2$ and $i+3$. Therefore, 1 now additionally approves of the winner on issue $i+1$ but does not approve the winners of issues $i+2$ and $i+3$ any more. Hence, free-riding led to a lower satisfaction for the free-riding voter. □

The same holds for the following large class of sequential OWA rules, as well as the sequential egalitarian rule.

Proposition 6. *Consider a sequential α-OWA rule such that there exists an $n \geq 8$ for which α^n is nonincreasing and satisfies $\alpha_3 > \alpha_{n-2}$. Then, free-riding can reduce the satisfaction of the free-riding voter.*

Proposition 7. *Free-riding can decrease the satisfaction of the free-riding voter under the sequential egalitarian rule.*

4 COMPUTATIONAL COMPLEXITY

In this section, we will study the computational complexity of free-riding. Overall, we will show that it is generally hard to do so. The reason for computational hardness, however, is a different one for optimization-based rules and for sequential rules. Observe that, due to the performance of, e.g., modern SAT- or ILP-solvers, computational hardness (in particular NP-completeness) cannot be seen as an unbreakable defense against manipulation. However, the main appeal of free-riding is its simplicity. A manipulator that is able to solve computationally hard problems has no benefit from restricting the potential manipulation to free-riding.

4.1 Free-Riding in Optimization-Based Rules

In this section, we study the computational complexity of free-riding for optimization-based rules. As our goal is to show that free-riding is hard, we start from a more fundamental problem: outcome determination. Indeed, any hypothetical free-rider needs to decide if, by voting dishonestly, the outcome would be better than the "truthful" outcome. To do so, she must be able to determine the outcome of an election. If this step turns out to be intractable, then we already have a computational barrier against free-riding. Hence, we study the following problem:

$\mathcal{R}$-Outcome Determination
Input: An election $\mathcal{E} = (N, \bar{A}, \bar{C})$, an issue i and a candidate $c \in C_i$.
Question: Does c win in issue i under $\mathcal{R}$?

In the following, we assume that for all f-Thiele rules, f is polytime computable.[6] Similarly, we assume that, for a given α-OWA rule and n voters, we can retrieve α^n in polynomial time. Now, we show that outcome determination is hard for both families of rules.

Theorem 8. *$\mathcal{R}$-Outcome Determination is NP-hard for every f-Thiele rule distinct from the utilitarian rule.*

Theorem 9. *$\mathcal{R}$-Outcome Determination is NP-hard for every α-OWA rule such that, for all n, α^n is nonincreasing and $\alpha_1 > \alpha_n$.*

Proof (Sketch). Fix a rule $\mathcal{R}$ satisfying the condition of the theorem. We show hardness by a reduction from CubicVertexCover, a variant of VertexCover where every node has a degree of exactly three [2]. Consider an instance (G, k) of this problem. Here, $G = (V, E)$ is a graph with n nodes and m edges where each node has a degree of exactly three, and $k \in \mathbb{N}$. We assume w.l.o.g. that $k < n$. We construct an instance of $\mathcal{R}$-Outcome Determination with $(k+1)$ issues and $3m$ voters. As $\alpha_1 > \alpha_{3m}$, there are two cases:

(1) There is a $p \in [2m]$ such that $\alpha_p > \alpha_{p+1}$, or
(2) There is a $p > 2m$ with $p < 3m$ such that $\alpha_1 = \alpha_p > \alpha_{p+1}$.

We sketch the proof for the first case. The full proof can be found in the full version of the paper. We construct an instance $(\mathcal{E}, k+1, c_{d_1})$ of $\mathcal{R}$-Outcome Determination. Here, we have one voter v_e for each edge $e \in E$, and two sets of dummy voters, $\{d_1, \ldots, d_p\}$ and $\{w_1, \ldots, w_{2m-p}\}$. In the first k issues, there is one candidate c_η for each node $\eta \in V$, plus one dummy candidate c_{d_i} for each dummy voter d_i. Here, each edge-voter v_e approves of the two node-candidates v_η and $v_{\eta'}$ such that $e = \{\eta, \eta'\}$. Moreover, each dummy voter d_i approves only of dummy candidate c_{d_i}, and all dummy candidates w_i approve of all candidates. In the last issue, there is one candidate c_v for all voters $v \in N \setminus \{w_i\}_{i \in [2m-p]}$, and every such v only approves of c_v. Finally, here, all voters in $\{w_i\}_{i \in [2m-p]}$ approve of all candidates.

The tie-breaking is defined as follows. We assume that each issue i is associated with a total ordering $\succ_i$ such that:

(1) If $i \in \{1, \ldots, k\}$, then node-candidates are preferred over other candidates, and $c_{d_n} \succ_i \cdots \succ_i c_{d_1}$;
(2) If $i = k+1$, then all candidates c_{v_e} (with $e \in E$) are preferred over other candidates, and $c_{d_1} \succ_i \cdots \succ_i c_{d_n}$;

We compare outcomes $\bar{w}$ and $\bar{w}'$ lexicographically, starting with issue 1. We want to show that (G, k) is a yes-instance if and only if $(\mathcal{E}, k+1, c_{d_1})$ is. Suppose that there exists a vertex cover for G with size at most k. Then, it can be shown that all edge-voters must win at least one issue in $[k]$, as increasing the satisfaction of a voter from 0 to 1 increases the OWA score more than increasing the

[6] This is justified by the fact that all relevant values of f can be computed ahead of time and stored in a look-up table.

satisfaction of a voter that has already positive satisfaction and edge voters are preferred in the tie-breaking. Let us show that c_{d_1} wins in $k+1$ if all edge-voters win at least once in issue in $[k]$. If voter d_1 never won an issue in $[k]$, then it means she has a satisfaction of 0. Since all edge-voters and all the w_i won at least once, there are at least $m + 2m - p = 3m - p$ voters with a satisfaction of at least 1. Therefore, d_1 occupies a position within the first p entries of the satisfaction vector, whereas all edge-voters occupy a position within the last $4m-p$ entries. Since $\alpha_p > \alpha_{p+1}$, in this case choosing in issue $k+1$ candidate c_{d_1} will yield a greater score than choosing a voter-candidate c_{v_e} for any edge $e \in E$. Finally, since c_{d_1} dominates in the tie-breaking every other candidate c_{d_j} in issue $k+1$, here we must choose c_{d_1}. On the other hand, suppose that d_1 wins at least one issue $i \in [k]$. Suppose – towards a contradiction – that c_{d_1} is not selected in issue $k+1$. Let c_v (for some voter $v \in N \setminus \{w_i\}_{i \in [2m-p]}$ distinct from d_1) be the candidate winning issue $k+1$. Observe that if we make c_{d_1} win in issue $k+1$ and make some candidate approved by v win in issue i, we would obtain a score that is higher or equal than before, and this would surely be preferred by tie-breaking: contradiction. We conclude that c_{d_1} must win in the final issue.

Now, suppose that there exists no vertex cover for G with size at most k. Then, there is one edge-voter that never wins an issue in $[k]$ (otherwise, some vertex cover would exist). By tie-breaking, this edge-voter would decide the last issue, i.e., c_{d_1} would not win. □

In light of this, one could conclude that free-riding is unfeasible for optimization-based rules. Still, one could argue – especially since we use worst-case complexity analysis – that sometimes the fact that a certain candidate wins can still be known (or guessed). For example, when a candidate receives an extremely disproportionate support, or when some external source (i.e., a polling agency having the computational power to solve $\mathcal{R}$-Outcome Determination) communicates the projected winners. In this case, the manipulator would need to solve a slightly different, potentially easier, problem: *Given that some candidate that I approve of wins in this specific issue, can I deviate from my honest approval ballot, without making this candidate lose*? Motivated by this, we study the following problem:

$\mathcal{R}$-Free-Riding Recognition

Input: An election $\mathcal{E} = (N, \bar{A}, \bar{C})$, an issue i, a candidate $c \in C_i$ such that $c \in \mathcal{R}(\mathcal{E})_i$, and a voter v such that $c \in A_i(v)$.

Question: Can v free-ride in $\mathcal{E}$ on issue i?

We define Generalized $\mathcal{R}$-Free-Riding Recognition analogously. Luckily, the picture does not change: this problem is still computationally hard for essentially the same families of rules.

Theorem 10. *(Generalized) $\mathcal{R}$-Free-Riding Recognition is* NP-*hard for every f-Thiele rule distinct from the utilitarian rule.*

Theorem 11. *(Generalized) $\mathcal{R}$-Free-Riding Recognition is* NP-*hard for every α-OWA rule for which there is a $c \geq 3$ such that, for every $n \in \mathbb{N}$, there is a nonincreasing vector α of size ℓ (with $3n \leq \ell \leq cn$) such that $\alpha_1 > \alpha_\ell$ and $\alpha_{3n} > 0$.*

Proof. We show hardness by a reduction from CubicVertexCover. Consider an instance (G, k) of this problem. Here, $G = (V, E)$ is a graph with n nodes and m edges where each node has a degree of exactly three, and $k \in \mathbb{N}$. By the condition of the theorem, we know there is an $\ell \geq 3m$ (polynomial in the size of m) such that $\alpha = (\alpha_1, \ldots, \alpha_\ell)$ contains at least $3m$ non-zero entries and $\alpha_1 > \alpha_\ell$. We will construct an instance of $\mathcal{R}$-Free-Riding Recognition with $(k+1)$ issues and ℓ voters. Since $\alpha_1 > \alpha_\ell$, we can distinguish essentially the same two cases as in the proof of Theorem 9. We treat here the first case. The second case and the treatment of the generalized problem are similar, and a full proof can be found in the full version of the paper.

We construct an instance $(\mathcal{E}, k+1, v_{e^*}, c_{d_1})$ of $\mathcal{R}$-Free-Riding Recognition (here, $e^* \in E$ is some edge, it does not matter which). The construction is similar to the one shown in the first case of the proof of Theorem 9. However, here, in issue $k+1$ voter v_{e^*} approves only of c_{d_1}, and we have $\ell - m - p$ dummy voters w_i instead of $3m - p$. The latter change makes no difference in our construction.

First, note that $(\mathcal{E}, k+1, v_{e^*}, c_{d_1})$ is indeed a legal instance of $\mathcal{R}$-Free-Riding Recognition, as surely c_{d_1} wins in issue $k+1$. If (G, k) is a yes-instance then we have already shown that this candidate wins, and here it is only receiving increased support. If it is a no-instance, then c_{d_1} will be supported by one voter that never won in the first k issues (namely, d_1), as well as by v_{e^*}. Since $\alpha_{p+m} \geq \alpha_{3m} > 0$ and since the edge-voters together with the dummy voters d_i occupy at most the first $p+m$ positions of the satisfaction vector, v_{e^*} will break the tie in favour of c_{d_1}.

Now, if (G, k) is a yes-instance of CubicVertexCover, then v_{e^*} can free-ride in the last issue: if she votes for her voter-candidate, then we have an election identical to the one constructed in the first case of the proof of Theorem 9, and we have already shown there that c_{d_1} wins if (G, k) has a vertex cover.

If (G, k) is a no-instance, then there are two cases: either v_{e^*} won in some issue in $[k]$ or not. If she did, there will be at least one voter v_e (with $e \in E \setminus \{e^*\}$) that never did, whose voter-candidate will get at least the same score as c_{d_1} (since v_{e^*} does not approve of the latter when she free-rides): c_{d_1} cannot win here. If she did not, there are again two cases: either v_{e^*} approves of some dummy candidate c_{d_i} (with $i > 1$) or of some c_{v_e} (where $e \in E$). In the first case, c_{d_i} would get a strictly higher score than c_{d_1}, while in the second case c_{v_e} would get at least the same score as c_{d_1} (and win by tie-breaking). In all cases, c_{d_1} loses: no free-riding is possible. □

Theorem 12. *(Generalized) $\mathcal{R}$-Free-Riding Recognition is* coNP-*hard for every α-OWA rule for which there is a $c \geq 2$ such that, for every $n \in \mathbb{N}$, there is a nonincreasing vector α of size ℓ (with $n < \ell \leq cn$) such that $\alpha_1 > \alpha_\ell$ and $\alpha_{\ell-n+1} = 0$.*

Theorems 10, 11 and 12 strengthen our previous observations. We conclude that free-riding is generally unfeasible for optimization-based rules, since the manipulator cannot even decide efficiently whether free-riding is possible. Next, we tackle sequential rules.

4.2 Free-Riding in Sequential Rules

In this section, we study the complexity of free-riding for sequential rules. First of all, observe that the computational barriers we exhibited in the previous section are not applicable here. Indeed, the outcome of a sequential rule is always poly-time computable: for every round, we can just iterate over all the candidates involved in that issue and pick the one maximizing the score up to that point.

However, although voters can easily verify if free-riding is possible, it might be still hard to judge its long-term consequences. If this is unfeasible, voters might be discouraged from free-riding (as it can have negative consequences). Hence, we study the following:

$\mathcal{R}$-Free-Riding
Input: An election $\mathcal{E} = (N, \bar{A}, \bar{C})$ and a voter $v \in N$.
Question: Is there an election $\mathcal{E}^*$ such that v can free-ride in $\mathcal{E}$ via $\mathcal{E}^*$ and $sat_{\mathcal{E}}(v, \mathcal{R}(\mathcal{E})) < sat_{\mathcal{E}}(v, \mathcal{R}(\mathcal{E}^*))$?

The problem of Generalized $\mathcal{R}$-Free-Riding is defined analogously. Now, we show that free-riding is NP-complete for a large class of sequential f-Thiele rules and the egalitarian rule.

Theorem 13. *$\mathcal{R}$-Free-Riding is NP-complete for every sequential f-Thiele rule for which there exists a $\ell \in \mathbb{N}$ such that (i) for all $j, j' \in [\ell]$ it holds $f(j) = f(j')$ and (ii) f is strictly decreasing on $\mathbb{N} \setminus [\ell - 1]$.*

The conditions of Theorem 13 apply to all functions that are constant up to a certain number ℓ, and from ℓ on become strictly decreasing. This is the case, e.g., for the sequential PAV rule.

Theorem 14. *$\mathcal{R}$-Free-Riding is NP-complete for the sequential egalitarian rule.*

Proof (Sketch). Membership is clear. To show hardness, we reduce from 3-SAT [22] and sketch the proof of its correctness. The full proof can be found in the full version of the paper.

Let ϕ be a 3-CNF with n variables and m clauses. We assume w.l.o.g. that ϕ is not satisfied by setting all variables to false and that each clause C_j contains exactly three literals. We construct an instance of $\mathcal{R}$-Free-Riding with $2(n + 1)$ voters and $5n + 1$ rounds. In particular, we will have two voters v_i and $\bar{v}_i$ for each variable x_i, a voter u, and a distinguished voter v, the manipulator.

In all rounds except for $5n + 1$, there are two candidates, c and $\bar{c}$. We assume that c always loses in ties (also in the final round). We group the first $4n$ rounds into n quadruples, e.g., quadruple 1 consists of rounds $(1, 2, 3, 4)$. In the first round of any such quadruple i, all voters approve of $\bar{c}$. In the second round of i, voters v and $\bar{v}_i$ vote for c, while voters u and v_i vote for $\bar{c}$; everyone else approves of both. In the third round of i, voters v and u approve of c and $\bar{c}$, respectively, and everyone else approves of both. In the final round of i, voter v votes for both c and $\bar{c}$, while everyone else votes for c. Next, in all rounds from $4n + 1$ to $5n - 1$, v approves of $\bar{c}$, u of both candidates, and everyone else of c. In round $5n$, v votes for $\bar{c}$, whereas every one else votes for c. Finally, in round $5n + 1$, there are $m + 1$ candidates, namely $c, c_1, \ldots, c_j$. Here, u approves of all candidates, voter v_i (resp. $\bar{v}_i$) approves of c and of all candidates c_j such that $x_i \notin C_j$ (resp. $\bar{x}_i \notin C_j$). Finally, voter v approves only of c.

We show correctness as follows. First, we observe that v can free-ride only in the first round of every quadruple: everywhere else, either she is losing, or her vote changes the outcome. Secondly, in all quadruples, if v votes truthfully in the first round, c and $\bar{c}$ win the second and third rounds, respectively; if she free-rides, the opposite happens. Note that, regardless of whether v free-rides or not, she will be satisfied with three rounds per quadruple (w.r.t. her honest preferences), and she will win as many rounds as u. Next, let ℓ be the (calculated) satisfaction of v and u after round $4n$. We can show that, for all pairs of voters v_i and $\bar{v}_i$, one voter won $s := \ell + n - 1$ rounds, while the other won $s + 1$ rounds, depending on whether c or $\bar{c}$ wins in the second round of quadruple i. Then, one can show that in all rounds from $4n + 1$ to $5n - 1$ only v and u win, and only v wins in round $5n$. Hence, before round $5n + 1$, all voters have satisfaction either s (including u) or $s + 1$ (including v).

We can interpret c winning in the second round of quadruple i as setting x_i to true. Crucially, voter v_i (resp. $\bar{v}_i$) has satisfaction $s + 1$ if $\bar{c}$ (resp. c) wins there, and s otherwise. We claim that c wins round $5n + 1$ iff this assignment satisfies ϕ. Briefly, if a clause C_j contains at least one satisfied literal, the minimal satisfaction if c_j wins would be s; otherwise, it would be $s + 1$. Since c also gives a minimal satisfaction of $s + 1$ and loses all ties, our claim follows.

Finally, observe that the (true) satisfaction of v from the first $5n$ rounds is exactly $4n$, irrespectively of whether she free-rides or not. If v never free-rides, $\bar{c}$ wins in the second round of every quadruple, and since we assumed that ϕ is not satisfied by setting all variables to false, v loses the last round. Hence, v can only raise her satisfaction to $4n + 1$ by winning the last round. To do so, she needs to force a satisfying assignment for ϕ by free-riding. It follows that ϕ is satisfiable if and only if v can manipulate via free-riding. ◻

We now consider the weaker notion of generalized free-riding. With this, we prove NP-completeness for a broader class of rules.

Theorem 15. *Generalized $\mathcal{R}$-Free-Riding is NP-complete for every sequential f-Thiele rule distinct from the utilitarian rule such that $f(i) > 0$ holds for every $i \in \mathbb{N}$.*

Theorem 16. *Generalized $\mathcal{R}$-Free-Riding is NP-complete for every sequential α-OWA rule such that, for all n, $\alpha = (\alpha_1, \ldots, \alpha_n)$ is nonincreasing and $\alpha_1 > \alpha_n$.*

5 NUMERICAL SIMULATIONS

So far, we have seen that sequential Thiele and OWA rules are generally susceptible to free-riding. However, we have seen that free-riding can be detrimental to the free-rider, i.e., her satisfaction can decrease. In this section, we use numerical simulations to shed more light on the *risk of free-riding* with sequential rules.

We use the following setup. We assume that voters and candidates are points in a 2-dimensional space; this is known as the 2d-Euclidean model [10, 18, 19, 23]. We sample both candidates and voters from a uniform distribution on a unit square. Voters' points are the same for all issues, candidates are sampled separately for each issue. A voter approves the closest candidate as well as any candidate that is similarly close (within +20% distance). We consider multi-issue elections with $n = 20$ voters, $k = 20$ issues, and 4 candidates per issue. Our results are based on 1000 elections.

In our experiments, we consider a subclass of Thiele methods and a subclass of OWA rules. For better comparison, we parameterize both classes with a parameter x (albeit this parameter has a different interpretation in both classes). We consider f-Thiele rules with $f_x(i) = i^{-x}$ for $x \in \{0, 0.25, 0.5, \ldots\}$. Note that for $x = 0$ this is the utilitarian rule, for $x = 1$ it is PAV, and for increasing x it approaches the leximin rule. Further, we consider α-OWA rules with

$$\alpha_x = (\underbrace{1, \ldots, 1}_{n-x \text{ many}}, \frac{1}{kn}, \frac{1}{k^2 n^2}, \ldots) \qquad \text{for } x \in \{0, 1, 2, \ldots\}.$$

Note that also this class contains the utilitarian rule ($x = 0$) and the leximin rule ($x = n - 1$).

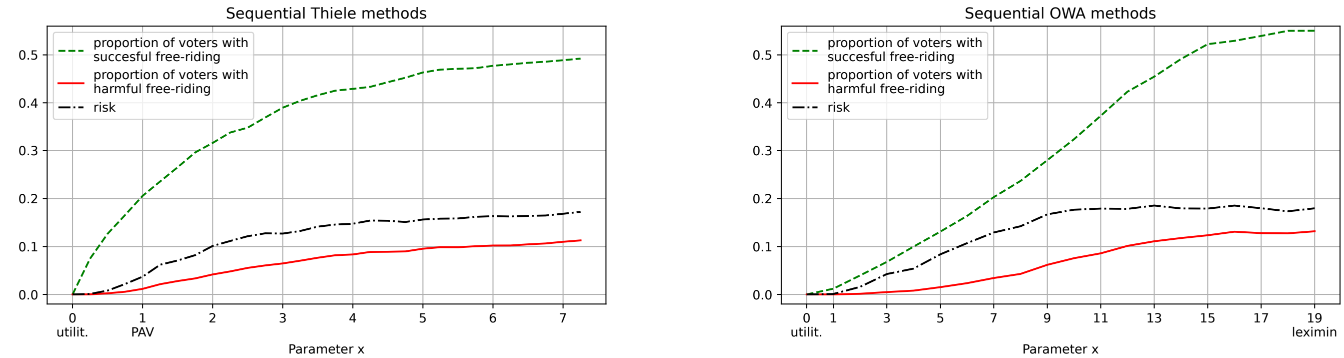

Figure 1: Results of the numerical simulations.

Within this model, we answer three questions: (Q1) How many voters have the possibility to increase their satisfaction by free-riding? (Q2) For how many voters can free-riding lead to a worse outcome? (Q3) What is the average risk of free-riding? Let us make these three questions precise. For each multi-issue election, we iterate over all voters and all issues and check whether free-riding is possible (Definition 1). That is, we only consider free-riding in single issues (and not repeated free-riding in more than one issue). Note that for a fixed occurrence of free-riding (i.e., in a specific issue, by a specific voter) it is computationally easy to determine the outcome when using sequential Thiele or sequential OWA rules.

Given an election, a voter, and an issue, we speak of *successful free-riding* if the voter can free-ride and this increases her satisfaction; we speak of *harmful free-riding* if the voter can free-ride but this decreases her satisfaction. Note that free-riding can also be neutral (with no change in satisfaction).

Figure 1 shows our results. We answer Q1 by displaying the proportion of voters with the possibility of successful free-riding (in at least one issue), averaged over all elections. Analogously, Q2 corresponds to the proportion of voters with the possibility of harmful free-riding (in at least one issue), averaged over all elections. We note that voters can have both the possibility of successful and harmful free-riding (on separate issues). Finally, for Q3, we define the risk of a voter in an election as the number of issues where harmful free-riding occurs divided by the number of issues where either successful or harmful free-riding occurs. Figure 1 shows the risk averaged over all voters (for whom successful or harmful free-riding is possible) and over all elections.

Let us discuss Figure 1. We clearly see that rules closer to the utilitarian rule ($x = 0$) are less susceptible to free-riding than those closer to leximin (larger values of x). We also see that – as expected – the utilitarian rule is the only rule where free-riding is not possible (cf. Proposition 2). We note that this increase in susceptibility (with distance to the utilitarian rule) has also been observed by Barrot et al. [6] for arbitrary manipulations. Both the proportion of voters that can successfully free-ride and those with the possibility of harmful free-riding grow with parameter x. The most important conclusion from this experiment is that the risk of free-riding is considerable (3.7% for sequential PAV, 17.2% for sequential leximin). This shows that harmful free-riding is not merely a theoretical possibility, but might be a phenomenon that indeed decreases the temptation of free-riding.

Finally, we briefly describe the impact of our chosen model parameters. Increasing the number of voters decreases the chance of voters being pivotal. Consequently, we would see a decrease in both successful and harmful free-riding. For a larger number of voters, it would make sense to move to a model where groups of voters free-ride. This requires additional assumptions about voter coordination (cf. the framework of iterative voting [35]). Varying the number of candidates leads to comparable results. Increasing the number of issues significantly increases the possibility of both successful and harmful free-riding, as effects may materialize only in the long run. In general, further simulations indicate that the general pervasiveness of harmful free-riding does not depend on our chosen parameter values.

6 DISCUSSION AND RESEARCH DIRECTIONS

We have seen that free-riding is an essentially unavoidable phenomenon in multi-issue voting (Theorem 3). However, we have also shown that there are computational issues to overcome for voters that would like to assess the consequences of free-riding. In particular for sequential voting rules, we have observed the possibility of negative outcomes for free-riders. Numerical simulations show that the frequency of harmful free-riding is non-negligible. This led us to the conclusion that it is less obvious how and when to free-ride than it seems at first sight. Another detriment to free-riding comes from the social context. In small groups, it may be obvious to other group members that free-riding takes place and thus can entail negative social consequences. Consequently, free-riding in real-world applications of multi-issue decision making may be less relevant than the theoretical possibility would suggest.

We conclude this paper with specific technical open problems. First, we would like to point out that many of our hardness proofs use several candidates per issue. Do all of these results still hold for binary elections? Second, our classification of sequential OWA rules with potentially harmful free-riding is not complete. Are there sequential OWA rules where free-riding is never harmful except for the utilitarian rule? Third, all our results apply to resolute rules, i.e., rules returning exactly one outcome. This condition can be lifted by introducing set extensions for comparing sets of outcomes (as done by Barrot et al. [6]). Would this change affect our conclusions? Finally, there are further voting rules to be considered, such as rules based on Phragmén's ideas [12, 37].

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their feedback. Martin Lackner was supported by the Austrian Science Fund (FWF), grant P31890. Jan Maly was supported by the FWF, grant J4581. Oliviero Nardi was supported by the Vienna Science and Technology Fund (WWTF), grant ICT19-065, the FWF, grant P32830, and the European Union's Horizon 2020 research and innovation programme, grant 101034440.

REFERENCES

[1] David S. Ahn and Santiago Oliveros. 2012. Combinatorial Voting. *Econometrica* 80, 1 (2012), 89–141.

[2] Paola Alimonti and Viggo Kann. 2000. Some APX-Completeness Results for Cubic Graphs. *Theoretical Computer Science* 237, 1 (2000), 123–134.

[3] Georgios Amanatidis, Nathanaël Barrot, Jérôme Lang, Evangelos Markakis, and Bernard Ries. 2015. Multiple Referenda and Multiwinner Elections Using Hamming Distances: Complexity and Manipulability. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015).* 715–723.

[4] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[5] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015).* 107–115.

[6] Nathanaël Barrot, Jérôme Lang, and Makoto Yokoo. 2017. Manipulation of Hamming-Based Approval Voting for Multiple Referenda and Committee Elections. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017).* 597–605.

[7] Sirin Botan. 2021. Manipulability of Thiele Methods on Party-List Profiles. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2021).* 223–231.

[8] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. 1997. Voting on Referenda: The Separability Problem and Possible Solutions. *Electoral Studies* 16, 3 (1997), 359–377.

[9] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. 1998. The Paradox of Multiple Elections. *Social Choice and Welfare* 15, 2 (1998), 211–236.

[10] Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2019. An Experimental View on Committees Providing Justified Representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019).* 109–115.

[11] Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2020. Electing Successive Committees: Complexity and Algorithms. In *Proceedings of the 34th Conference on Artificial Intelligence (AAAI-2020)*, Vol. 34. 1846–1853.

[12] Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. 2017. Phragmén's Voting Methods and Justified Representation. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017).* AAAI Press, 406–413.

[13] Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. 2022. Approval-Based Apportionment. *Mathematical Programming* (2022), 1–29.

[14] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. 2021. Justified Representation for Perpetual Voting. *IEEE Access* 9 (2021), 96598–96612.

[15] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair Public Decision Making. In *Proceedings of the 2017 ACM Conference on Economics and Computation.* 629–646.

[16] Vincent Conitzer and Lirong Xia. 2012. Paradoxes of Multiple Elections: An Approximation Approach.. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-2012).*

[17] Théo Delemazure, Tom Demeulemeester, Manuel Eberl, Jonas Israel, and Patrick Lederer. 2022. Strategyproofness and Proportionality in Party-Approval Multiwinner Elections. *arXiv preprint arXiv:2211.13567* (2022).

[18] Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. What Do Multiwinner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017).* 494–501.

[19] James M. Enelow and Melvin J. Hinich. 1990. *Advances in the Spatial Theory of Voting.* Cambridge University Press.

[20] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 2, 27–47.

[21] Eric J. Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. 2019. Fair and Efficient Memory Sharing: Confronting Free Riders. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI-2019)*, Vol. 33. 1965–1972.

[22] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractibility: A Guide to NP-Completeness.* Freeman.

[23] Michał T. Godziszewski, Paweł Batko, Piotr Skowron, and Piotr Faliszewski. 2021. An Analysis of Approval-Based Committee Rules for 2D-Euclidean Elections. In *Proceedings of the 35th Conference on Artificial Intelligence (AAAI-2021).* 5448–5455.

[24] Theodore Groves and John Ledyard. 1977. Optimal Allocation of Public Goods: A Solution to the "Free Rider" Problem. *Econometrica: Journal of the Econometric Society* (1977), 783–809.

[25] Boas Kluiving, Adriaan de Vries, Pepijn Vrijbergen, Arthur Boixel, and Ulle Endriss. 2020. Analysing Irresolute Multiwinner Voting Rules with Approval Ballots via SAT Solving. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-2020).* 131–138.

[26] Martin Lackner. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th Conference on Artificial Intelligence (AAAI-2020).* AAAI Press.

[27] Martin Lackner and Jan Maly. 2021. Perpetual Voting: The Axiomatic Lens. In *Proceedings of the 8th International Workshop on Computational Social Choice (COMSOC 2021).*

[28] Martin Lackner and Jan Maly. 2023. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th Conference on Artificial Intelligence (AAAI-2023).* to appear.

[29] Martin Lackner and Piotr Skowron. 2018. Approval-Based Multi-Winner Rules and Strategic Voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018).* ijcai.org, 340–436.

[30] Martin Lackner and Piotr Skowron. 2021. Consistent Approval-Based Multi-Winner Rules. *Journal of Economic Theory* 192 (2021), 105173.

[31] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences.* Springer. https://doi.org/10.1007/978-3-031-09016-5

[32] Jérôme Lang and Lirong Xia. 2009. Sequential Composition of Voting Rules in Multi-Issue Domains. *Mathematical Social Sciences* 57, 3 (2009), 304–324.

[33] Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press.

[34] Barton E. Lee. 2018. Representing the Insincere: Strategically Robust Proportional Representation. *arXiv preprint arXiv:1801.09346* (2018).

[35] Reshef Meir. 2017. Iterative Voting. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, 69–86.

[36] Dominik Peters. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018).* 1549–1557.

[37] Lars Edvard Phragmén. 1895. *Proportionella val. En valteknisk studie.* Lars Hökersbergs förlag, Stockholm.

[38] Paul A. Samuelson. 1954. The Pure Theory of Public Expenditure. *The Review of Economics and Statistics* (1954), 387–389.

[39] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús A. Fisteus, Pablo Basanta Val, and Piotr Skowron. 2017. Proportional Justified Representation. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-2017).* AAAI Press, 670–676.

[40] Luis Sánchez-Fernández and Jesús A. Fisteus. 2019. Monotonicity Axioms in Approval-Based Multi-Winner Voting Rules. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2019).* 485–493.

[41] Markus Schulze. 2004. Free Riding. *Voting Matters* (2004), 2–8. Issue 18.

[42] Piotr Skowron, Piotr Faliszewski, and Jérôme Lang. 2016. Finding a Collective Set of Items: From Proportional Multirepresentation to Group Recommendation. *Artificial Intelligence* 241 (2016), 191–216.

[43] Piotr Skowron and Adrian Górecki. 2022. Proportional Public Decisions. In *Proceedings of the 36th Conference on Artificial Intelligence (AAAI-2022).* 5191–5198.

[44] Thorvald N. Thiele. 1895. Om Flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger.* 415–441.

[45] Ronald R. Yager. 1988. On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decisionmaking. *IEEE Transactions on systems, Man, and Cybernetics* 18, 1 (1988), 183–190.