Proportionality Guarantees in Elections with Interdependent Issues

Markus Brill, Evangelos Markakis, Georgios Papasotiropoulos, Jannik Peters

Abstract

We consider a multi-issue election setting over a set of possibly interdependent issues with the goal of achieving proportional representation of the views of the electorate. To this end, we employ a proportionality criterion suggested by Skowron and Górecki [2022], that guarantees fair representation for all groups of voters of sufficient size. For this criterion, there exist rules that perform well in the case where all the issues have a binary domain and are independent of each other. In particular, this has been shown for *Proportional Approval Voting* (PAV) and for the *Method of Equal Shares* (MES). In this paper, we go two steps further: we generalize these guarantees for issues with a non-binary domain, and, most importantly, we consider extensions to elections with dependencies among issues, where we identify restrictions that lead to analogous results. To achieve this, we define appropriate generalizations of PAV and MES to handle conditional ballots. In addition to proportionality considerations, we also examine the computational properties of the conditional version of MES. Our findings indicate that the conditional case poses additional challenges and differs significantly from the unconditional one, both in terms of proportionality guarantees and computational complexity.

1 Introduction

Proportional representation of voters' preferences is an important desideratum in a wide range of social choice settings, including parliamentary elections [20], approval-based multiwinner voting [13], and participatory budgeting [19]. In two recent papers, Freeman et al. [10] and Skowron and Górecki [23] studied a *public decision* model, where a collective decision needs to be taken on a given set of unrelated binary (yes/no) issues. In real life scenarios, these kinds of situations are often handled via majority rule: an issue is implemented if and only if more than 50% of the voters are in favor of it; and while this might be the most straightforward decision rule, it comes at a disadvantage: It is not proportional. For instance, consider a scenario where three binary decisions need to be taken, and the electorate is split into two groups. The first group of voters, which makes up two thirds of the electorate, votes "yes" on every issue, while the second group votes "no" on all issues. Majority rule would set all issues to "yes," completely ignoring the preferences of one third of the electorate. One might argue that setting two issues to "yes" and one issue to "no" would be a more fair solution. Such considerations led Freeman et al. [10] and Skowron and Górecki [23] to adopt several proportionality notions and proportional voting rules from the setting of approval-based multiwinner voting to their model.

The model studied by the above-mentioned previous works has two important limitations. Firstly, the issues under consideration are only allowed to be binary, which makes it impossible to model scenarios where an issue has more than two possible alternatives. Secondly, their setting does not allow any dependencies between the issues, although in some situations it might be the case that the approval of an issue is dependent on the decision taken for another. The more general setting for proportional decision-making that we study here is based on the *conditional approval voting* framework introduced by Barrot and Lang [5]. This setting allows for more than two choices per issue as well as for dependencies between the issues; consequently, it addresses both limitations discussed above.

Example 1. *Consider a scenario, where a municipality has decided that the following projects will be funded: (1) a public park, (2) a pedestrian infrastructure, and (3) a community center. However, the location of these sites is not fixed yet, and the voters should decide on whether to build each project at the Southside, the Centralside, or the Northside district. Obviously, the domain of each issue here is of size three. Voter 1, a resident of Southside, is voting in favor of building the park in her neighborhood and therefore casts an (unconditional) approval ballot for this option. The preferences of the other voters are more complicated: Voter 2 is concerned about the traffic congestion that a new community center could cause and votes for building a community center at any location only if a pedestrian infrastructure is built at the same location as well. Finally, voter 3 doesn't want any two projects built in the same location. Hence her approval for any project depends on the decisions made for the other two projects.*

In this paper, we focus on the proportionality concept that was introduced by Skowron and Górecki [23], which provides strong guarantees for all groups of voters of sufficient size. These guarantees are considered as more powerful than other previously studied ones, which were only able to ensure proportional representation to groups of voters that have similar preferences. Such "cohesive" groups are rare to be found in practice in the unconditional setting [6], let alone when dependencies between issues exist.

1.1 Contribution.

The main challenge left open by Skowron and Górecki [23] lies in incorporating in their model dependencies between issues, which can be seen as a stepping-stone to the building of a theory of fairness for general public decision settings. Motivated by this, we first introduce appropriate generalizations of *Proportional Approval Voting* (PAV) and the *Method of Equal Shares* (MES), which are two prominent rules with provable guarantees on proportional representation for binary, unrelated issues. We then make progress along two fronts. First, we generalize the known guarantees to elections over issues with non-binary domains. Secondly, and most importantly, we consider elections with dependencies among issues, where we identify sufficient restrictions that lead to analogous results. Our work is the first that provides guarantees of proportionality for elections with conditional ballots. Moreover, for the conditional version of MES, we also study computational aspects. Our results demonstrate that the conditional setting poses additional challenges and differs significantly from the unconditional one, both in terms of proportionality and complexity. As an example, we show that MES is hard to implement in the general conditional case, whereas it has been known to be tractable in unconditional elections. Another highlight of our work is that PAV and MES achieve proportionality bounds under different assumptions (that have some degree of complementarity), and thus, one cannot reach yet an absolute conclusion when compared against each other, under the conditional framework, in contrast to the unconditional setting.

1.2 Related Work.

On a high level, there are three lines of research that are related to our paper. First, our work is closely related to approval-based multiwinner voting [13]. Starting with the works of Aziz et al. [3] and Sánchez-Fernández et al. [21], a plethora of proportionality notions have been defined and studied. Originally defined for multiwinner voting, proportionality notions have more recently been extended to more general settings such as participatory budgeting [19, 7]. Quantitative proportionality notions have also been studied by Aziz et al. [4] and Skowron [22].

Secondly, our work is related to multi-issue decision-making and seat-posted elections. With the exception of the aforementioned papers by Skowron and Górecki [23] and Freeman et al. [10], these have not been studied from a proportionality perspective. Instead, the literature has focused on issues such as Anscombe's and Ostrogorski's paradoxes [1, 11, 14] or axiomatic comparisons between preferences and decisions over issues and the entire decision space [9, 2].

Third, our work is utilizing a model of conditional approval voting over combinatorial domains. This was introduced by Barrot and Lang [5], who also proposed three voting rules for incorporating dependencies between issues and studied (mainly axiomatic) properties. Later on, the works by Markakis and Papasotiropoulos [16, 17], focused on the algorithmic aspects of the winner determination problem under the conditional analog of the (minisum) approval voting rule and also considered the computational complexity of controlling the election outcome, under the same rule. Apart from approval-based elections, the presence of preferential dependencies remains a major challenge and several frameworks have been considered, as extensively discussed by Lang and Xia [15] and Chevaleyre et al. [8].

2 Preliminaries

In this section, we introduce the setting of conditional approval voting, we define the proportionality notion we are interested in, and define the two voting rules that will be the focus of our work.

2.1 Conditional Election Setting

In describing conditional approval elections, we closely follow the notation and terminology of the previous papers on the topic, by Barrot and Lang [5] and Markakis and Papasotiropoulos [16, 17]. We consider a group $N = \{v_1, v_2, \ldots, v_n\}$ of n *voters* that needs to make a decision over a set of m possibly interdependent *issues* $I = \{I_1, I_2, \ldots, I_m\}$. Each issue I_j, is associated with a *domain* $D_j = \{a_j^1, a_j^2, \ldots, a_j^d\}$ that corresponds to the alternatives for this issue. Without loss of generality, we assume that $|D_j| = |D_{j'}| \geq 2, \forall j, j' \in [m]$ (if not, one can add dummy alternatives) and we denote the domain size by d.

A *conditional approval voting instance*, or simply an *instance*, is determined by a tuple $P = (I, N, B)$, where B denotes the conditional approval ballots of the set of voters N over the issues of I. In what follows, we describe the format of conditional approval ballots.

Each voter $v_i \in N$ is associated with a directed graph, whose vertex set coincides with the set of issues, referred to as her *dependency graph*, $G_i = (I, E_i)$. A directed edge $(I_k, I_j) \in E_i$, means that according to voter v_i, issue I_j is affected by I_k. We use $\Gamma_i(I_j)$ to denote the (possibly empty) set of direct predecessors, i.e., in-neighbors, of issue I_j in G_i and $\Gamma_i^*(I_j) = \Gamma_i(I_j) \cup \{I_j\}$. Voters cast conditional approval ballots that are expressed as follows: For an issue I_j with $|\Gamma_i(I_j)| = 0$, voter v_i casts a standard (unconditional) approval ballot, stating explicitly all the alternatives of D_j that are approved by her, the number of which varies from 0 to d. For the case that $|\Gamma_i(I_j)| > 0$, voter v_i needs to specify all the combinations of alternatives for issues in $\Gamma_i^*(I_j)$ that she approves, i.e., that make her satisfied w.r.t. issue I_j. These combinations are expressed in the form $\{s : t\}$, where $s \in \times_{I_\ell \in \Gamma_i(I_j)} D_\ell$, and $t \subseteq D_j$. Such a ballot signifies the satisfaction (i.e., approval) of a voter with respect to issue I_j, when the in-neighbors of I_j in G_i are set to the alternatives specified by s, and the selection for I_j belongs to t. We note that we do not impose that voters submit a ballot for every issue (they can abstain if they are not satisfied with any outcome w.r.t. certain issues). In the case that all voters have the same dependency graph, we let $\Delta_{\texttt{in}}$ denote the maximum in-degree of this graph.

Given a conditional approval voting instance, the *global dependency graph* is the undirected simple graph G that results from ignoring the orientation of edges in the graph $(I, \bigcup_{i\in[n]} E_i)$, where E_i is the edge set of v_i's dependency graph. We use $\Gamma(I_j)$ to denote the (possibly empty) set of neighbors of issue I_j in G and $\Gamma^*(I_j) = \Gamma(I_j) \cup \{I_j\}$. Moreover, Δ denotes the maximum degree of any vertex in G. Note that if $\Gamma(I_j) = \emptyset$ for all $j \in [m]$, then a conditional approval voting election degenerates to a classical approval election over m issues.

For any $I' \subseteq I$, any tuple $w \in \times_{I_j \in I'} D_j$, i.e., that includes an alternative from D_j for every issue $I_j \in I'$, is referred to as a *suboutcome*. If w specifies a value for all issues of an instance $P = (I, N, B)$, we simply call it an *outcome* of P. A *conditional approval voting rule* is a function that maps each conditional approval voting instance P to an outcome w. Given an outcome w, we let $u_i(w)$ denote the number of issues with respect to which voter v_i is satisfied under w. If w is clear from the context, we simply refer to this as u_i. The following example illustrates our setting.

Example 2. *Consider an instance where $n = 3$ voters decide upon $m = 3$ issues, namely I_1, I_2, I_3, with domain size $d = 3$ and domains $D_i = \{x_i, y_i, z_i\}$, $i \in \{1, 2, 3\}$. The conditional approval ballots of the voters are specified in the table that follows, in which, an entry "–" indicates that the voter cannot be satisfied w.r.t. the corresponding issue under any outcome.*

	voter v_1	voter v_2	voter v_3
issue I_1	x_1	–	x_1, y_1
issue I_2	x_2, y_2, z_2	$x_3 : x_2$ $y_3 : y_2$ $z_3 : z_2$	y_2
issue I_3	–	x_3	$y_1 y_2 : (x_3, z_3)$

The dependency graphs of the voters, namely G_1, G_2, G_3 are depicted in the figure below, together with the global dependency graph G, which has maximum degree $\Delta = 2$.

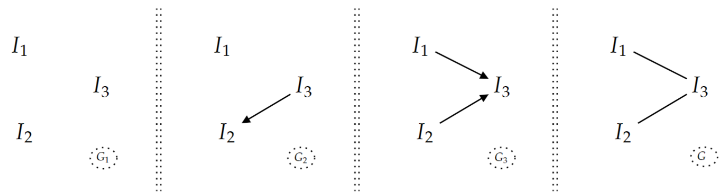

2.2 Proportionality Criterion

We now define a (parameterized) notion of proportionality that generalizes the one suggested by Skowron and Górecki [23]. The basic rationale behind any proportionality definition is the intuitive idea that any fraction of voters should have the ability to influence a corresponding fraction of decisions:

> *A group of voters that makes up a β-fraction of the electorate should be able to decide on βm issues.*

In particular, such a criterion requires that minority opinions are represented as well (proportionally to their size).

To formalize this notion, an important parameter is the set of issues that a voter can be possibly satisfied with. Namely, for a voter $v_i \in N$, we let R_i denote the set of issues

$I_j \in I$, for which v_i has submitted at least one conditional approval ballot (or unconditional if $\Gamma_i(I_j) = \emptyset$), i.e., there is at least one selection of alternatives for $\Gamma_i^*(I_j)$ that makes v_i satisfied w.r.t. issue I_j. Formally, $R_i = \{I_j \in I : \exists w \in \times_{I_\ell \in \Gamma_i^*(I_j)} D_\ell$ that satisfies v_i w.r.t. $I_j\}$. To avoid trivialities, we assume that $|R_i| > 0$ for all $v_i \in N$. By definition, voter v_i cannot be satisfied with respect to any issue in $I \setminus R_i$, under any outcome. In Example 2, we have $R_1 = \{I_1, I_2\}$, $R_2 = \{I_2, I_3\}, R_3 = \{I_1, I_2, I_3\}$.

For a group $V \subseteq N$ of voters, we let r_V denote the number of issues for which all voters of V approve at least one alternative, i.e., $r_V = | \cap_{i \in V} R_i|$. In Example 2, we have $r_N = 1, r_{\{v_1,v_2\}} = 1$ and $r_{\{v_2,v_3\}} = 2$. The role of r_V is important in the definition below in which, the proportionality guarantee of a group of voters takes into account the maximal number of issues that all members of the group care about.

Definition 1. *A conditional voting rule is α-proportional, for some $\alpha \in [0, 1]$, if for every conditional approval voting instance $P = (I, N, B)$ with $|N| = n$ and for every $V \subseteq N$, there exists a voter $v_i \in V$ such that if w is the winning outcome under the considered rule, then*

$$u_i(w) > \alpha r_V \frac{|V|}{n} - 1.$$

The parameter α in the definition represents the degree of proportionality that a voting rule can guarantee. Ideally, we would like to have α-proportional rules for $\alpha = 1$, as this would mean that the elected outcome aligns with the views of the electorate in a proportional manner. However, as we will soon show, such a rule does not exist and more relaxed values will need to be considered. In the unconditional case, and with a binary domain for each issue, it was shown by Skowron and Górecki [23] that $\alpha = \frac{1}{2}$ is achievable. In our more general setting, we will see that the degree of proportionality cannot be expressed by a constant; rather, it will be a function of the input instance, dependent on d and Δ.

2.3 Conditional Voting Rules

We focus on two conditional rules, that constitute natural generalizations of their well-studied unconditional versions.

Conditional Proportional Approval Voting (cPAV)

An outcome w of a conditional approval election in an instance $P = (I, N, B)$ gains a score of $\sum_{k=1}^{u_i(w)} \frac{1}{k}$ from every voter v_i who is satisfied with w with respect to $u_i(w)$ issues. The cPAV score of w is $\sum_{v_i \in N} \sum_{k=1}^{u_i(w)} \frac{1}{k}$, or, in words, it is the sum of the scores that it gains from all the voters of the electorate. The outcome that achieves the highest cPAV score is the winning one under cPAV. Note that the only difference between the unconditional definition of PAV [24, 12] and the version we suggest here, comes from the way that a voter's satisfaction w.r.t. an issue is defined.

Conditional Method of Equal Shares (cMES)

The unconditional version of cMES was introduced by Peters and Skowron [18] (and was originally being referred to as "Rule X"). In our more general setting, the rule consists of two phases, and the first one works in rounds. Initially, each voter is given a budget m, equal to the number of issues and all issues are being considered as unfixed. Fixing an issue will cost a price of n. Unlike in the unconditional version of the rule, we also allow that several issues are fixed at the same time, in which case the price to be paid is n times the number of issues to be fixed. For any round t of the first phase and an issue $I_j \in I$ that has not been fixed yet, denote by $\Gamma^u(I_j)$ the set of all issues in $\cup_{v_i \in N} \Gamma_i^*(I_j)$ that remain unfixed until round t. For every such issue I_j, for every set $I' \subseteq \Gamma^u(I_j)$, and for every possible (sub)outcome w on the issues of I', we perform the following: First, we identify the set of voters $S(w)$ who have a positive remaining budget and are satisfied with respect to I_j under w; second, we calculate[1] the price $p(w)$, which is such that if each voter in $S(w)$ paid $p(w)$ or all the money she has left, then the voters from $S(w)$ would altogether pay $n \cdot |I'|$. Finally, among the above, we determine the set of issues I' and the suboutcome w with a minimal value for $p(w)$; we reduce the budget of every voter in $S(w)$ by $p(w)$ (or to 0 if their current budget is less than $p(w)$); we set the decision on the issues of I' to w, and we continue with the next round, until no further purchase can be made. It might happen that after this procedure, there are issues for which the decision has not been set. For these, in the second phase, we select an alternative arbitrarily.

A natural case for elections with conditional ballots, that we pay special attention to, is when all the voters agree on the dependencies among issues, i.e., when they have the same dependency graph. In some scenarios, this may even be enforced by the election organizer, either for uniformity reasons or when there are obvious enough dependencies among issues that apply to all voters. With a common graph, the execution of cMES becomes a bit simpler. In particular, for a yet unfixed issue $I_j \in I$, we only look at the single subset of the in-neighborhood of I_j, say I', that has not been fixed in the previous rounds. For this set I', we check all possible suboutcomes w to identify the voters who can get satisfied with respect to I_j, and continue in the same manner as in the description of cMES above.

3 Conditional PAV

We begin our study by examining cPAV, the conditional version of Proportional Approval Voting; a voting rule that exhibits significant proportionality guarantees in the binary and unconditional case [23]. The main result of this section is the identification of a proportionality guarantee under a certain assumption (Theorem 1). Before delving into this, we present an example that establishes that, as in the unconditional setting, we should not have too high expectations in terms of the α-value that is achievable.

Example 3. *Consider an instance $P = (I, N, B)$, with m issues, such that m is a multiple of $d^{\Delta+1}$ and furthermore, suppose that m can be written as $m = k(\Delta + 1)$, for an integer k. Assume also that the issues of I can be partitioned in k sets, namely $M_1, M_2, \dots, M_k$, of size $(\Delta + 1)$ each, so that issues from M_i are not dependent on issues from M_j for any voter, and for any $i \neq j$. Let there also be $d^{\Delta+1}$ voters. We fix an $i \in [k]$ and we focus on a single set of issues M_i. We will make each issue of M_i dependent on all the remaining Δ issues of the group M_i, for all voters. This means that the global dependency graph is a disjoint union of cliques.*

To describe the voters' preferences, note that there are exactly $d^{\Delta+1}$ possible outcomes for the issues of any single clique. We define the preferences so that for every possible outcome in each clique, there is exactly one voter who is satisfied with that outcome, with respect to all the $\Delta + 1$ issues. Furthermore this voter is dissatisfied in any other outcome with respect to all these issues. For instance, say we fix a suboutcome $x = (x_1, \dots, x_{\Delta+1})$ for a particular clique. Then, we will have exactly one voter whose ballot is $\{x_{-1} : x_1,\ x_{-2} : x_2,\ \dots\ , x_{-(\Delta+1)} : x_{(\Delta+1)}\}$, where x_{-i} denotes the tuple $(x_1, \dots, x_{i-1}, x_{i+1}, \dots x_{\Delta+1})$. Thus, observe that if x is indeed the final selection for the clique under consideration, the satisfaction of the corresponding voter, with respect to these issues, equals exactly $\Delta + 1$. We use the same construction of preferences for the issues of the remaining cliques. Hence, for every $V \subseteq N$, it holds that $r_V = m$.

[1] The calculation of $p(w)$ for a given w can be done easily, see e.g. the proof of Theorem 4.

Observation 1. *There does not exist a voting rule that is α-proportional for $\alpha = 1$.*

Proof. Consider the instance P of Example 3, pick an arbitrary voting rule and let w be its winning outcome in P. Since the voters do not agree on any issue under any outcome, there exists at least one voter v_i satisfied with at most $\frac{m}{n}$ issues. If v_i is satisfied with strictly less than $\frac{m}{d^{\Delta+1}}$ issues, then consider the set $V = \{v_i\}$. Proportionality requires that $u_i(w) \geq m\frac{|V|}{n} = \frac{m}{d^{\Delta+1}}$, which is a contradiction. Hence, it must be that all voters are satisfied with exactly $\frac{m}{d^{\Delta+1}}$ issues. If we focus now on $V = N$, it should hold that $u_i(w) > m - 1$ for all $v_i \in N$, which is not the case. □

In fact, Example 3 has further negative implications.

Remark 1. *Consider a rule that is reasonably fair to the voters, in the sense that it does not treat any voter in a significantly different manner. Then in Example 3, every voter would have to be satisfied with respect to exactly $\frac{m}{d^{\Delta+1}}$ issues. But then, looking at $V = N$ would imply that the rule cannot be α-proportional, for any $\alpha > \frac{1}{d^{\Delta+1}}$.*

The previous discussion highlights that most probably, the best we could expect is a proportionality guarantee of $\alpha = \frac{1}{d^{\Delta+1}}$. We have not yet been able to obtain such a general result for all instances. The main technical difficulty with conditional elections is the analysis of the PAV score. This is more challenging to handle now because the satisfaction of a voter can change more abruptly when we alter the value of a single issue, due to the effect this may have on other issues. On the positive side, we can prove a guarantee by restricting the voters' preferences. It remains an open problem to determine whether such restrictions could be relaxed or removed.

Assumption 1. *For every voter $v_i \in N$, any two issues in R_i are either located in different components or are at a distance of at least 4 from each other in the global dependency graph.*

In other words, the assumption states that if a voter can be satisfied w.r.t. some issue, she cannot be satisfied w.r.t. to any other "nearby" issue in the global dependency graph. Hence it intends to capture voters who express preferences for a rather limited number of issues per component. The most instructive, and intuitively simplest, yet non-trivial, case in which the assumption holds, appears when voters express preferences only for one issue per connected component. As a first example, consider instances with a common dependency graph for all voters, consisting of a collection of in-stars, where the voters may only be satisfied with respect to the central issue of each star, dependent on the leaf-issues. Generalizing this, we can also allow voters having different dependency graphs, as long as, again, there is a central issue in each component of the global dependency graph that the voters care about, albeit expressing different dependencies for it. Furthermore each voter could specify a different issue that she cares about in each component, justified by the voter's specialization, expertise or prioritization.

Theorem 1. *Under Assumption 1, cPAV is α-proportional for $\alpha = \frac{1}{(1+\Delta^2)d^{\Delta+1}}$.*

Proof. Consider a conditional election instance (I, N, B) in which Assumption 1 holds, let G be its global dependency graph, and w its cPAV winning outcome. For any fixed issue I_j, our proof will be based on a counting argument, where we will examine all possible ways of changing the alternatives chosen for I_j and its neighbors in w. Let s_j be a (sub)outcome that specifies an alternative for every issue in $\Gamma^*(I_j)$. Let $\delta_w(s_j)$, be the difference of the cPAV score if we change every issue of $\Gamma^*(I_j)$ from its value in w to the value indicated by s_j, minus the cPAV score of w. We denote by $\mathcal{F}_j$ the set of all possible suboutcomes s_j such that $s_j(I') \neq w(I')$, for at least one issue $I' \in \Gamma^*(I_j)$, so that s_j does impose an outcome different from w. Given that $|\Gamma^*(I_j)| \leq \Delta + 1$, it holds that $|\mathcal{F}_j| \leq d^{\Delta+1} - 1$.

The main component of the proof is the estimation of the expression $\sum_{I_j \in I} \sum_{s_j \in \mathcal{F}_j} \delta_w(s_j)$, denoted by $\mathcal{S}_w$ for simplicity. Before we proceed, recall also that if u_i is the number of issues that v_i is satisfied with under w, the voter contributes to the cPAV score the quantity $1 + \frac{1}{2} + \cdots + \frac{1}{u_i}$. At what follows we will use $\Gamma(\Gamma(I_j))$ to denote the set of issues that are at distance 2 from I_j in G. We consider now two cases.

First, we consider an issue I_k with respect to which voter v_i is satisfied under w. We will take into account all possible tuples s_j in the expression $\mathcal{S}_w$, that may affect the satisfaction of v_i for this issue. These are precisely the suboutcomes $s_k \in \mathcal{F}_k$, the suboutcomes $s_\ell \in \mathcal{F}_\ell$ that correspond to any issue $I_\ell \in \Gamma(I_k)$, and also the suboutcomes $s_\ell \in \mathcal{F}_\ell$ for any issue $I_\ell \in \Gamma(\Gamma(I_k))$. Let us denote by $\mathcal{F}'_k$ the set of all these suboutcomes. In worst case, any change of w by the values indicated by some $s \in \mathcal{F}'_k$ may cause voter v_i to become dissatisfied with respect to issue I_k. Observe also that by Assumption 1, the suboutcomes from $\mathcal{F}'_k$ cannot affect any other issues that voter v_i is satisfied with under w, in other words, v_i will remain satisfied with all the other $u_i - 1$ issues that she was satisfied with under w. Therefore, when we change w according to a suboutcome $s \in \mathcal{F}'_k$, either the voter v_i does not contribute anything to $\mathcal{S}_w$, or her contribution will be negative, and equal to $-1/u_i$ due to the definition of cPAV. To bound the number of tuples that can create this negative contribution, we need to estimate $|\mathcal{F}'_k|$. To do this, note that the set $\Gamma^*(I_k) \cup \Gamma(\Gamma(I_k))$ contains exactly I_k itself, its neighbors, and its neighbors of neighbors which altogether correspond to at most $(1 + \Delta + \Delta(\Delta - 1)) = 1 + \Delta^2$ issues. Given also that for every such issue I_j, the number of different tuples $s_j \in \mathcal{F}_j$ to consider is $d^{\Delta+1} - 1$, it is true that $|\mathcal{F}'_k| \leq (1 + \Delta^2)(d^{\Delta+1} - 1)$. Thus, the total contribution of voter v_i is in worst case at least $-\frac{(1+\Delta^2)(d^{\Delta+1}-1)}{u_i}$. Since we know that under w, there are exactly u_i issues that v_i is satisfied with, we conclude that for every $i \in [n]$ the contribution of v_i to $\mathcal{S}_w$ is at most $-\frac{u_i(1+\Delta^2)(d^{\Delta+1}-1)}{u_i}$.

Second, consider the $|R_i| - u_i$ issues that v_i is dissatisfied under w, but for which we know that there exists an assignment of values that satisfy v_i (by definition of R_i). Fix any such issue, say I_k. We know that in $\mathcal{F}_k$ there exists at least one suboutcome, say s_k, that can make v_i satisfied with respect to I_k. By Assumption 1, we know that changing w according to s_k cannot affect the remaining issues that v_i is satisfied with under w. Hence, there is at least one suboutcome, s_k, for which voter v_i contributes a positive value to $\mathcal{S}_w$, which equals $\frac{1}{u_i+1}$. Therefore, for every $i \in [n]$ the contribution of v_i to $\mathcal{S}_w$ is at least $\frac{|R_i|-u_i}{u_i+1}$.

The optimality of w implies that $\delta_w(s_j) \leq 0$, for every $s_j \in \mathcal{F}_j$ and for every $j \in [m]$. Furthermore, $-(d^{\Delta+1} - 1) \leq 0$ and hence the following hold:

$$0 \geq \sum_{I_j \in I} \sum_{s_j \in \mathcal{F}_j} \delta_w(s_j) \geq \sum_{v_i \in N : u_i > 0} -\frac{u_i(1+\Delta^2)(d^{\Delta+1}-1)}{u_i} + \sum_{v_i \in N} \frac{|R_i| - u_i}{u_i + 1} \geq$$
$$\sum_{v_i \in N} \left(-(1+\Delta^2)(d^{\Delta+1} - 1) + \frac{|R_i| - u_i}{u_i + 1} \right) = \sum_{v_i \in N} \left(\frac{|R_i| + 1}{u_i + 1} - (1+\Delta^2) d^{\Delta+1} \right)$$

Equivalently, $\sum_{v_i \in N} \frac{|R_i|+1}{u_i+1} \leq (1 + \Delta^2)d^{\Delta+1}n$, and if we fix any set $V \subseteq N$, it holds that $\sum_{v_i \in V} \frac{|R_i|+1}{u_i+1} \leq (1 + \Delta^2)d^{\Delta+1}n$. Using the fact that $|R_i| \geq r_V$, and by a rearrangement of the terms, we also have that

$$\sum_{v_i \in V} \frac{1}{u_i + 1} \leq \frac{(1+\Delta^2)d^{\Delta+1}n}{r_V + 1}. \qquad (1)$$

Due to the harmonic and arithmetic mean inequality, $\frac{n}{\sum_{i\in[n]}\frac{1}{x_i}} \leq \frac{1}{n}\sum_{i\in[n]} x_i$, we get

$$\sum_{v_i\in V}\frac{1}{u_i+1} \geq \frac{|V|^2}{\sum_{v_i\in V}(u_i+1)} = \frac{|V|^2}{|V|+\sum_{v_i\in V} u_i}. \tag{2}$$

Combining the relations (1) and (2) gives (after reordering)

$$\frac{1}{|V|}\sum_{v_i\in V} u_i \geq \frac{|V|}{n}\frac{r_V+1}{(1+\Delta^2)d^{\Delta+1}} - 1 > \frac{1}{(1+\Delta^2)d^{\Delta+1}}\frac{|V|}{n}r_V - 1.$$

Consequently, for every $V \subseteq N$, there exists a voter $v_i \in V$, for which $u_i(w) > \frac{1}{(1+\Delta^2)d^{\Delta+1}}\frac{|V|}{n}r_V - 1$. □

Despite the negative results indicated by Example 3, Theorem 1 provides the first guarantee of proportionality for conditional approval elections. On the downside, even in the unconditional setting, determining the winning outcome for PAV is NP-hard.

Implications for the Unconditional Case. Assumption 1 trivially holds if all voters submit unconditional ballots since then, any two issues belong to different components in the global dependency graph (which does not have any edges). For this case, $\Delta = 0$ and the α-value achieved by Theorem 1 is $\frac{1}{d}$. This strictly generalizes the result of Skowron and Górecki [23], that deals only with the case of $d = 2$, and had left as an open problem the cases with higher values of d. Note also that Example 3 still works for $\Delta = 0$, which makes the result of Theorem 1 tight in the unconditional case, for any domain size, generalizing once again the analogous result.

4 Conditional MES

The computational intractability of PAV in the unconditional setting motivated the study of other rules, such as the Method of Equal Shares (MES), that overcome computational barriers and at the same time have desirable characteristics from the perspective of proportional representation. In this section, we extend this line of work to conditional elections, by focusing on the conditional version of MES as defined in Section 2.3. We start with studying algorithmic properties of cMES, and then proceed to proportionality guarantees.

4.1 Computational Complexity of cMES

The main result of this subsection is that Conditional MES is not, in general, computable in polynomial time, which is a noteworthy characteristic that differentiates it from its unconditional variant. This holds even for binary domains and rather simple dependency graphs, when either the maximum in-degree is large (Theorem 2) or the dependency graphs of the voters do not coincide (Theorem 3). Despite these negative results, there are well-motivated restricted families of instances for which we can compute cMES in polynomial time. In particular, when neither of the aforementioned conditions hold, a winning outcome can be computed efficiently (Theorem 4). Finally, for the case of different dependency graphs, we have also identified a restriction that implies polynomial-time computability (Theorem 5).

Theorem 2. *The winning outcome under cMES cannot be computed in polynomial time, unless P=NP, even when the voters have a common dependency graph.*

In the proof of Theorem 2, we have used in an essential way the fact that the maximum in-degree of the common dependency graph is non-constant. However, we expect that this represents rather extreme cases, and that in practical scenarios for conditional voting, it is more important to focus on the case of bounded in-degree, which has also been the focus of previous works on computational aspects of such elections, see e.g. [17]. Even in this case, the NP-hardness remains, when the voters have different dependency graphs.

Theorem 3. *Assuming ties are broken in favor of the largest set of buyers, the winning outcome under cMES cannot be computed in polynomial time, unless P=NP, even with a constant maximum in-degree in each voter's dependency graph.*

Proof. To prove the statement, we reduce from 3SAT. Given a 3SAT instance, Π, on q variables, namely $x_1, x_2, \ldots, x_q$, and r clauses, namely $c_1, c_2, \ldots, c_r$, we create a conditional election instance $P = (I, N, B)$ of $|I| = q+r+1$ issues and $|N| = q+r+1$ voters, as follows:

- We create a binary issue I_0 of domain $\{z_1, z_2\}$ and we add a voter voting unconditionally for z_1.
- For every $j \in [q]$, i.e. for every variable x_j of Π, we add a binary issue I_j of domain $\{t_j, f_j\}$. We refer to these issues as variable-issues. Furthermore, we add a variable-voter for every $j \in [q]$, who is only voting for $\{t_j : z_1, f_j : z_1\}$ and is dissatisfied with any other issue.
- For every clause c_j of Π, $j \in [r]$, we add a binary issue I'_j of domain $\{\texttt{pos}_j, \texttt{neg}_j\}$. We refer to these as clause-issues. Furthermore, we add a clause-voter who only cares to be satisfied w.r.t. I_0, and is voting for $\{c_j \wedge \texttt{pos}_j : z_1\}$, where c_j contains at most 3 variable-issues.

Observe that the maximum in-degree in every voter's dependency graph is at most 4. Furthermore, every voter is only interested in getting satisfied w.r.t. a single issue I_0.

Lemma 1. *The instance Π is a YES-instance if and only if there exists an outcome that satisfies all voters w.r.t. I_0.*

In the remaining proof we will show that one cannot efficiently determine the set of voters who should pay for the first purchase in a run of cMES, unless P=NP. To do this, we establish the following claim: the set of all voters N, is the largest in cardinality set of voters that can jointly buy a set of alternatives at a minimum per voter cost in the first iteration if and only if the 3SAT formula is satisfiable.

For the forward direction, suppose that the formula is satisfiable. Then by Lemma 1, there is an outcome that satisfies all voters. This implies that the voters can buy an outcome of all the $q + r + 1$ issues, and the per voter cost would be $p(w) = q + r + 1$. If this was not the minimum possible per voter cost in the first iteration of cMES, then there was a different purchase, say for a suboutcome w', with $p(w) < q + r + 1$, due to the tie-breaking rule. We need to check whether there exists such a set of κ voters, that are willing to buy a suboutcome w' on λ issues, where $\kappa < q + r + 1$, and such that $p(w') < q + r + 1$. But then, with λ issues, we have $p(w') = \frac{\lambda(q+r+1)}{\kappa}$ and $\frac{\lambda(q+r+1)}{\kappa} < q + r + 1$ if and only if $\lambda < \kappa$. However, it must be the case that $\lambda \geq \kappa$. To see why, consider an arbitrary set of voters S that includes a variable-voters, b clause-voters and possibly the voter who votes for I_0 unconditionally. In order for such a set of voters to buy a suboutcome, note that the issue I_0 has to be included, and furthermore, each variable-voter requires her own variable issue to be included as well. For every clause-voter, we also have to include the corresponding clause issue, hence in total the set of issues that will be fixed is at least $a + b + 1$. Thus, it is impossible that κ voters can decide to buy together a suboutcome on less than κ issues.

For the reverse direction, suppose that in the first iteration of cMES, the set selected to buy alternatives is the set of all voters. Since, the ballot of every voter contains a distinct issue, this means that the voters have bought an alternative for all the issues in order to be concurrently satisfied with I_0. Then, by Lemma 1, the SAT formula is satisfiable. Hence, we have shown that we can run efficiently the first iteration of cMES in the instance we constructed, if and only if we can decide if the SAT formula is satisfiable. □

Moving to positive results, we demonstrate below that when we have a common dependency graph and the in-degree is bounded, cMES can be implemented efficiently. We find this to be a natural case (which complements the tractability landscape for cMES with Theorems 2 and 3), since there are numerous scenarios where the voters are likely to have the same perception on the dependence structure among issues. Recall that in these instances, $\Delta_{\mathtt{in}}$ is the maximum in-degree.

Theorem 4. *If all the voters have the same dependency graph, and Δ_{in} is bounded by a constant, then the winning outcome of cMES can be computed in polynomial time.*

We conclude the computational analysis of cMES with one more positive result, applicable to the more general setting where voters can have different dependency graphs.

Theorem 5. *If each connected component of the global dependency graph has no more than a constant number of vertices, then the winning outcome under cMES can be computed in polynomial time.*

4.2 Proportionality Considerations of cMES

We first prove that in the conditional case, there is a family of instances for which cMES is strictly worse in terms of proportionality than cPAV.

Proposition 1. *If $\Delta \geq 1$, for any $\rho \in \mathbb{R}_{\geq 1}$, cMES is not $\frac{1}{\rho}$-proportional, even for instances that satisfy Assumption 1.*

Therefore, restrictions of a different flavor than Assumption 1 are necessary in order to end up with a bounded proportionality guarantee for cMES. To understand when to expect a good behavior from cMES, it is important to revisit first the unconditional case, where an assumption was also needed to achieve any proportionality bound. More precisely, the guarantee for MES by Skowron and Górecki [23] was established under the assumption that there are no abstainers, i.e., every voter approves at least one alternative from every issue, hence $r_V = m$. In the conditional setting, we will use a generalization of this statement, in the following form: no matter how the in-neighbors of an issue are set, a voter can still be satisfied with at least one alternative of that issue.

Assumption 2. *For every issue I_j and for every voter v_i, we assume that for every combination of values for the issues in $\Gamma_i(I_j)$, there is a choice for I_j that satisfies v_i with respect to I_j. When $\Gamma_i(I_j) = \emptyset$, we simply assume that v_i approves at least one alternative from the domain of I_j.*

We note that Assumption 2 is incomparable to the assumptions needed for the polynomial algorithms of Section 4.1 and that it implies $r_V = m$, for any set of voters V.

Theorem 6. *Under Assumption 2, cMES is α-proportional for $\alpha = \frac{1}{(\Delta+1)d^{\Delta+1}}$. When all voters have the same dependency graph, then the same bound holds with Δ replaced by Δ_{in}, the maximum in-degree over all issues.*

For the case of a common dependency graph for all voters, we show that we can also go a small step further by slightly relaxing Assumption 2, in a way that Theorem 6 still holds and at the same time no significant improvements on the proportionality bound would be possible. Before that, we give the following notation: consider an instance $P = (I, V, B)$, the voters of which have the same dependency graph. Then for every issue $I_j \in I$, let $\Gamma_{\mathtt{in}}(I_j)$ be the set of in-neighbors of issue I_j in the dependency graph and let $Z(I_j) = \{I_k \in I \text{ such that } I_k \in \Gamma_{\mathtt{in}}(I_j) \wedge I_j \in \Gamma_{\mathtt{in}}(I_k)\}$.

Assumption 3. *For every issue I_j and for every voter v_i, we assume that for every combination of values for the issues in $\Gamma_{in}(I_j) \setminus Z(I_j)$, there is a choice of values for the issues in $Z(I_j) \cup \{I_j\}$ that satisfies v_i with respect to I_j. When $\Gamma_{in}(I_j) = \emptyset$, we simply assume that v_i approves at least one value from the domain of I_j.*

It is easy to verify that, for the case of a common dependency graph, Theorem 6 still works under this weaker assumption. The reason is that the issues in $Z(I_j)$ cannot be fixed before the round where I_j gets fixed, as they depend on it. Hence, in the proof of Theorem 6, we only need to impose the condition of Assumption 2 for the issues in $\Gamma_{\mathtt{in}}(I_j) \setminus Z(I_j)$ (the ones that might have been fixed before the round where I_j is bought). Furthermore, we exhibit below that under Assumption 3, a proportionality guarantee that is significantly better than the result of Theorem 6 is impossible.

Observation 2. *For $\Delta_{in} \geq 1$ and $d \geq 2$, cMES is not $\frac{1}{d^{\Delta_{in}+1}}$-proportional, even for instances that satisfy Assumption 3 and even if all voters have the same dependency graph.*

Implications for the Unconditional Case. In the unconditional case, cMES coincides with MES and is computable in polynomial time. A corollary of Theorem 6 that concerns the unconditional case is the generalization of the proportionality guarantee for MES from the case of binary decisions to any domain size d (which was left as an open question by Skowron and Górecki [23]), while meeting the lower bound of $\frac{1}{d}$ from Example 3.

5 Conclusions

We studied generalizations of two well-known voting rules with proportionality guarantees, PAV and MES, to the setting of conditional approval elections, in which voters' exhibit preferential dependencies between the issues under consideration. Our main results establish that both cPAV and cMES can achieve proportionality bounds, under different assumptions. PAV seems to favor situations where the satisfaction score of a voter is somewhat restricted, whereas MES has a better behavior when voters are "easier" to please for every issue.

There are several questions for potential future work, we first note that we do not have yet a complete picture about the tightness of our bounds. Also it has been challenging to understand whether the assumptions used can be relaxed, and to what extent. The assumption on MES seems to be quite critical even in its unconditional variant; as for PAV, we are optimistic that relaxations might be possible, given also the fact that under unconditional ballots, PAV works well without any assumptions. One can also study the behavior of other rules under the conditional setting, or think of further ways to generalize MES, based on how a purchase is made and who participates in each purchase. In general, we believe that proportional representation in combinatorial domains is a fascinating area, worth further exploration.

Acknowledgments

The work of Markus Brill and Jannik Peters has been supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1 and the Graduiertenkolleg “Facets of Complexity” (GRK 2434). The work of Georgios Papasotiropoulos has been supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “3rd Call for H.F.R.I. PhD Fellowships” (Fellowship Number: 5163). The work of Evangelos Markakis has been supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the “First Call for H.F.R.I. Research Projects to support faculty members and researchers and the procurement of high-cost research equipment” grant (Project Number: HFRI-FM17-3512). The authors would like to thank Jérôme Lang for graciously hosting Markus Brill, Georgios Papasotiropoulos and Jannik Peters concurrently at Paris Dauphine University, where this work was initiated. The visit of Markus Brill has been supported by a visiting professorship from Paris Dauphine University and the visit of Georgios Papasotiropoulos has been supported by COST Action CA16228 - GAMENET.

References

[1] Gertrude Elizabeth Margaret Anscombe. On frustration of the majority by fulfilment of the majority’s will. *Analysis*, 36(4):161–168, 1976.

[2] Fatma Aslan, Hayrullah Dindar, and Jean Lainé. When are committees of Condorcet winners Condorcet winning committees? *Review of Economic Design*, 26(3):417–446, 2022.

[3] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 2(48):461–485, 2017.

[4] Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez-Fernández, and Piotr Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 902–909, 2018.

[5] Nathanaël Barrot and Jérôme Lang. Conditional and sequential approval voting on combinatorial domains. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 88–94, 2016.

[6] Robert Bredereck, Piotr Faliszewski, Rolf Niedermeier, and Andrzej Kaczmarczyk. An experimental view on committees providing justified representation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 109–115, 2019.

[7] Markus Brill, Stefan Forster, Martin Lackner, Jan Maly, and Jannik Peters. Proportionality in approval-based participatory budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023.

[8] Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. Preference handling in combinatorial domains: From AI to Social Choice. *AI Magazine*, 29(4):37–37, 2008.

[9] Hayrullah Dindar and Jean Lainé. Compromise in combinatorial vote. *Social Choice and Welfare*, pages 1–32, 2022.

[10] Rupert Freeman, Anson Kahng, and David M. Pennock. Proportionality in approval-based elections with a variable number of winners. In *Proceedings of the 29th International Conference on International Joint Conference on Artificial Intelligence (IJCAI)*, pages 132–138, 2021.

[11] Robin Fritsch and Roger Wattenhofer. The price of majority support. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 436–444, 2022.

[12] Svante Janson. Phragmén's and Thiele's election methods. Technical report, arXiv:1611.08826 [math.HO], 2016.

[13] Martin Lackner and Piotr Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2022.

[14] Gilbert Laffond and Jean Lainé. Single-switch preferences and the ostrogorski paradox. *Mathematical Social Sciences*, 52(1):49–66, 2006.

[15] Jérôme Lang and Lirong Xia. Voting in combinatorial domains. In *Handbook of Computational Social Choice*, pages 197–222. Cambridge University Press, 2016.

[16] Evangelos Markakis and Georgios Papasotiropoulos. Computational aspects of conditional minisum approval voting in elections with interdependent issues. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 304–310, 2020.

[17] Evangelos Markakis and Georgios Papasotiropoulos. Winner determination and strategic control in conditional approval voting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 342–348, 2021.

[18] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 793–794, 2020.

[19] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, pages 12726–12737, 2021.

[20] Friedrich Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications*. Springer, 2014.

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

[22] Piotr Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 820–840, 2021.

[23] Piotr Skowron and Adrian Górecki. Proportional public decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5191–5198, 2022.

[24] Thorvald Nicolai Thiele. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441, 1895.

Markus Brill
University of Warwick, Coventry, UK
TU Berlin, Berlin, Germany
Email: markus.brill@warwick.ac.uk

Evangelos Markakis
Input Output Global (IOG)
Athens University of Economics and Business, Athens, Greece
Email: markakis@aueb.gr

Georgios Papasotiropoulos
Athens University of Economics and Business, Athens, Greece
Email: gpapasotiropoulos@aueb.gr

Jannik Peters
TU Berlin, Berlin, Germany
Email: jannik.peters@tu-berlin.de