Proportional Participatory Budgeting with Cardinal Utilities

Dominik Peters, Grzegorz Pierczyński, Piotr Skowron

Abstract

We study voting rules for participatory budgeting, where a group of voters collectively decides which projects should be funded using a common budget. We allow the projects to have arbitrary costs, and the voters to have arbitrary additive valuations over the projects. We formulate two axioms that guarantee proportional representation to groups of voters with common interests. To the best of our knowledge, all known rules for participatory budgeting do not satisfy either of the two axioms; in addition we show that the most prominent proportional rule for committee elections, Proportional Approval Voting, cannot be adapted to arbitrary costs nor to additive valuations so that it would satisfy our axioms of proportionality. We construct a simple and attractive voting rule that satisfies one of our axioms (for arbitrary costs and arbitrary additive valuations), and that can be evaluated in polynomial time. We prove that our other stronger axiom is also satisfiable, though by a computationally more expensive and less natural voting rule.

1 Introduction

A growing list of cities now uses Participatory Budgeting (PB) to decide how to spend their budgets [Cabannes, 2004, Aziz and Shah, 2020]. Through a voting system, PB allows the residents of a city to decide which projects will be funded by the government. This increases civic involvement in government, by increasing the number of issues that are decided by democratic vote, and by allowing residents to submit their own project proposals.

To count the votes, most cities use a variant of a simple protocol: Each voter is allowed to vote for a certain number of project proposals. Then, the projects with the highest number of votes are funded, until the budget limit is reached. While simple and intuitive, this is a bad voting rule. To see this, consider Circleville, a fictional city divided into four districts. A map of the city is shown in Figure 1. The districts all have similar sizes, but Northside has the largest population. Suppose \$400k have been allocated to PB, and suppose that all the project proposals are of a local character (such as school renovations), so we can assume that residents only vote for projects that concern their own district. For example, every Northside resident will cast votes for projects A, B, C, and D, but no one else votes for these. Because Northside is the most populous district, the Northside projects will all receive the highest number of votes, and the voting rule described will spend the entire budget on Northside projects. The 280k residents of the other districts are left empty-handed.

To circumvent this obvious issue, many cities have opted to hold separate elections for each district. The budget is divided in advance between the districts (e.g., in proportion to their number of residents), each project is assigned to a district, and voters only vote in their local election. While this avoids the issue of spending the entire budget in Northside, this fix introduces many other problems. For example, projects on the boundary of two districts (such as A and P) need to be assigned to one of them. Residents of the other district may be in favor of the boundary project, but cannot vote for it. Thus boundary projects are less likely to be funded, even if they would be more valuable overall. Similarly, projects without a specific location that benefit the entire city cannot be handled. Also, interest groups that are not geographic in nature will be underserved; for instance, parents across the city might favor construction of a large playground (project C), but with separate district elections,

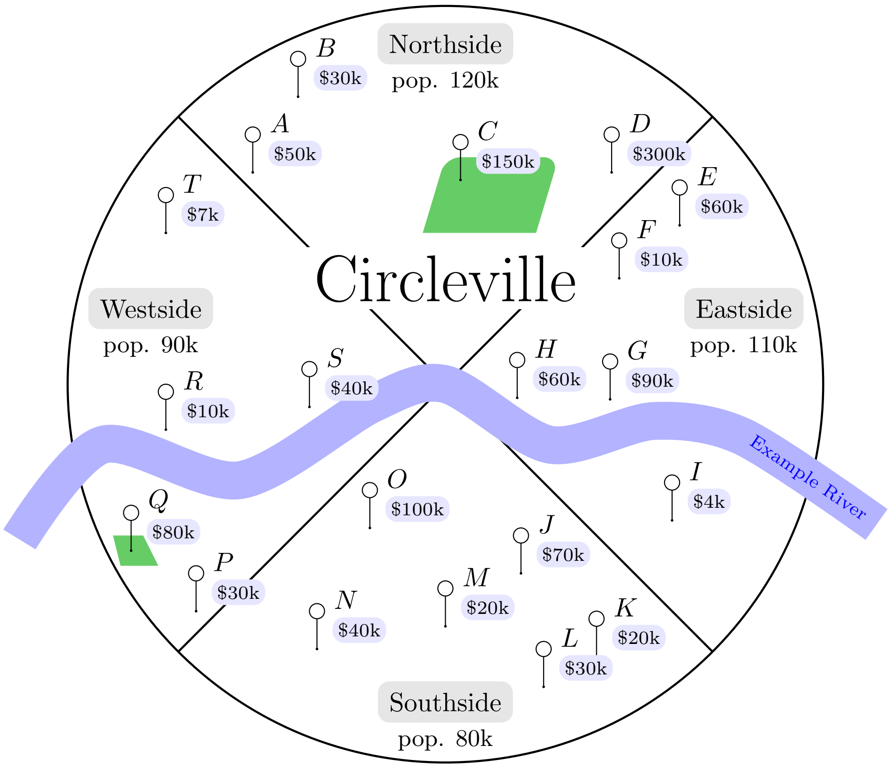

Figure 1: Map of Circleville, showing the locations and costs of the PB project proposals.

parents cannot form a voting block. Similarly, bike riders across the city cannot express their joint interest in the construction of a bike trail along Example River (projects R, S, H, and G).

To solve these problems, it seems desirable to hold a single city-wide election, but use a voting system that ensures that money is spent proportionally. The voting system should automatically and endogenously identify groups of voters who share common interests, and make sure that those groups are appropriately represented. This aim has been identified by several researchers [Aziz et al., 2018b], but no convincing proposal for a proportional voting rule has emerged so far. Indeed, no good formalization of "proportionality" for the PB context has been identified in the literature, except for the concept of the *core*. However, the core is a very demanding requirement, and there are situations where it fails to exist [Fain et al., 2018].

In this paper, we formalize proportionality for participatory budgeting as an axiom called *extended justified representation* (EJR). The axiom requires that no group of voters with common interests is underserved. We construct a simple and attractive voting rule that satisfies EJR for approval preferences, and that satisfies EJR up to one project for general additive valuations. We then discuss a potential strengthening of EJR, and show that this strengthening is still satisfiable, albeit by a different voting rule. We hope that our axioms and rules will provide a strong starting point for the further development of the PB literature from a social choice perspective.

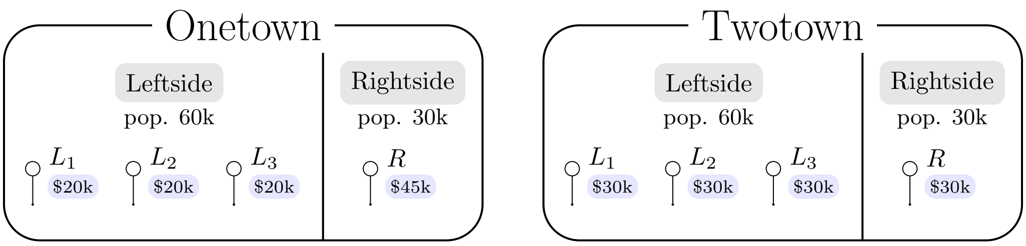

Figure 2: Onetown and Twotown are identical, except that the projects have different costs. Both have a budget of \$90k available for PB.

Our approach: Generalize concepts from multi-winner voting

Both our proportionality axiom and our voting rule are generalizations of concepts that have been introduced in the literature on multi-winner voting [Faliszewski et al., 2017]. That literature can be seen as handling a special case of PB, where all projects cost the same amount of money. This is often called the *unit cost assumption*. Under this assumption, the problem is equivalent to selecting a committee of a specified size k. It turns out that the unit cost assumption substantially simplifies the problem.

Much of the relevant literature studies rules that work with *approval ballots*, where voters are allowed to approve or disapprove each project. In the main part of the paper, we will allow any additive valuations (not just 0/1), which is more expressive. Indeed, the proportionality axioms and voting rules that we introduce all work for general additive valuations. This is notable, since allowing additive valuations introduces significant conceptual difficulty. Most prominent multi-winner voting rules seem to not naturally extend to additive valuations (or at least not gracefully). However, to compare our results to the literature, let us for now focus just on approval-based rules.

The study of approval-based multi-winner voting rules has been very productive [Aziz et al., 2017, Brill et al., 2017, Sánchez-Fernández et al., 2017, Lackner and Skowron, 2018, Peters and Skowron, 2020]. Researchers have identified a considerable number of proportionality axioms and of attractive voting rules for this case. We will begin our discussion by explaining why most of these results break down if we do not use the unit cost assumption.

Proportional Approval Voting is not proportional, nor is any variant of it

Probably the most popular multi-winner voting rule is Proportional Approval Voting (PAV), also known as Thiele's method after its Danish inventor Thiele [1895]. Thiele's rule is based on optimization. Suppose that N is the set of voters, and that each voter $i \in N$ has indicated a set $A(i)$ of projects that i approves. Then for each set W of projects which is feasible (i.e. its total cost is at most the budget limit), the rule computes the score

$$\text{PAV-score}(W) = \sum_{i \in N} \left(1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{|W \cap A(i)|}\right).$$

The output of PAV is a feasible set W that maximizes this score. When the unit cost assumption holds, PAV is a great rule and lives up to its name: it is known to be proportional both in an axiomatic sense [Aziz et al., 2017] and in a quantitative sense [Skowron, 2018]. In fact, among all optimization-based rules, only PAV is proportional [Lackner and Skowron, 2018, Aziz et al., 2017].

However, when the unit cost assumption does not hold, PAV ceases to guarantee proportional representation. To see this, consider the city of Onetown shown in Figure 2. Onetown has 90,000 residents split in two districts, and has \$90,000 available for participatory budgeting. The 60,000 residents of Leftside all vote for projects $\{L_1, L_2, L_3\}$ each of which costs \$20,000. The 30,000 residents of Rightside vote for the single project $\{R\}$ which costs \$45,000. Note that the prices are such that we can either afford to implement all three L-projects giving PAV score 110,000, or implement two L-projects and the R-project giving PAV score 120,000. Thus, PAV implements project R and only two L-projects. However, note that Leftside residents form two thirds of the population of Onetown, and so by proportionality are entitled to two thirds of the budget (\$60,000), which is enough to implement all three L-projects. Hence, Leftside is underrepresented by PAV.

To see what is going on, consider Twotown from Figure 2. Twotown is just like Onetown, except that now each projects cost \$30,000. Note that for Twotown it is still the case that we can either afford all three L-projects, or two L-projects plus R. By the same calculation as before, PAV implements the latter possibility. This time, this is the proportional choice: Leftside now deserves only two projects, since only two projects are affordable with Leftside's share of the budget.

Onetown and Twotown are nearly identical: same number of residents, same district structure, same alternatives, same approval sets, and even the feasibility constraint (three L or two L plus R) is the same. Since the definition of PAV only depends on these characteristics, it must select the same outcome for both towns. But the prices differ, and therefore different outcomes are proportional, and hence PAV fails proportionality. The same is true for all other rules that depend only on preferences and feasibility constraints but not prices. This suggests that there is no variant of PAV that retains its proportionality guarantees beyond the unit cost case.

Theorem 1. *Every voting rule that only depends on voters' utility functions and the collection of budget-feasible sets must fail proportionality, even on instances with a district structure.* □

(An instance has *district structure* if projects and voters can be partitioned into disjoint districts such that each voter approves exactly the projects belonging to the voter's district.)

Rule X: A simple method that guarantees proportionality

Peters and Skowron [2020] recently proposed a new committee voting rule called Rule X. It turns out that this rule can be naturally extended beyond unit costs, and that the resulting rule does not suffer from the defects in Thiele's method.

Rule X starts out by dividing the available budget into equal parts, and giving each voter their share. On a high level, Rule X then repeatedly looks for a project whose approvers have enough money left to fund it; it does so until no further projects are affordable. Notice that any rule of this type is proportional on instances with a district structure. For example, in Onetown, the residents of Leftside receive \$60k in total, so can afford L_1, L_2, and L_3. But to ensure good behavior on other instances it is crucial exactly how the rule chooses among different affordable projects and how the rule divides the chosen project's cost among its supporters. Rule X always spreads the cost of the project as *evenly as possible*, which means that all supporters contribute the same amount of money to it; if some supporters do not have this amount of money left, they spend their entire remaining budget. If several projects are affordable, Rule X chooses the project that minimzes the highest amount that any supporter needs to pay. (Thus, all else equal, Rule X favors cheap projects over expensive ones, and favors projects with many supporters over projects with fewer.)

As we mentioned, it is clear that Rule X is proportional on district-based instances. On its own, this is a rather weak guarantee. In the real world, like in Circleville (Figure 1), voters will sometimes vote for projects in other districts, and it is uncommon that voters will approve *all* the projects in their own district. A truly proportional voting rule should be able to represent all kinds of interest groups, even in cases where the same voter is part of several such groups.

Consider an arbitrary subset of voters, $S \subseteq N$. For example, S could be the residents of a district, or the set of parents in Circleville, or the set of bike users. The group S forms a fraction $|S|/|N|$ of the population, and thus intuitively its members deserve to control a fraction $|S|/|N|$ of the budget. This idea is the basis of most proportionality axioms developed in the literature; they differ by how they formalize the notion of "deserving" part of the budget. We will consider an axiom that guarantees to represent groups whenever they are sufficiently *cohesive*, in the sense of having similar preferences. Suppose that S can come up with a set T of projects such that T can be funded with a $|S|/|N|$ fraction of the budget. Suppose further that each voter in S approves all the projects in T; this means the group is cohesive. Then an axiom called Extended Justified Representation (EJR) demands that the voting rule select a set W such that at least one voter in S approves at least $|T|$ of the funded projects in W. In other words, EJR prohibits sets W where all the voters in S are underrepresented in the sense that they would all prefer the set T to W.[1]

EJR was first proposed for committee voting by Aziz et al. [2017]. EJR is a demanding property, with PAV one of the very few known voting rule satisfying it, but we have seen that without unit costs, PAV fails EJR even in well-structured cases. However, Peters and Skowron [2020] showed that Rule X also satisfies EJR, and one of our main results shows that this also holds without the unit cost assumption.

Theorem. *Rule X satisfies Extended Justified Representation.*

The intuition behind this result is that, under Rule X, a group S is explicitly given their share of the budget. As the rule progresses, the money of S is spent and by design of Rule X it is spent on projects that provide good value for money. Thus, the only way that S could end up underrepresented is if Rule X does not spend all of S's money; but we can show that this never happens if S is cohesive.

Among other known committee voting rules, we know of only one that can easily be extended beyond unit costs: Phragmén's rule [1894, 1895] which was proposed at the same time as Thiele's rule. Phragmén's rule satisfies some proportionality axioms in the unit cost case [Janson, 2016, Brill et al., 2017, Peters and Skowron, 2020], and Aziz et al. [2018b] show that, in contrast to Thiele's method, it can be naturally extended to cases without unit costs. The rule is easiest to explain via a continuous process. Each voter is assigned a virtual bank account, which starts out empty. We continuously top up each voter's account at a constant rate, say \$1 per hour. We continue doing this until the first moment where there exists a project whose supporters own enough money to finance that project. We then implement that project and reset the bank accounts of the supporters to 0. (If several projects become affordable simultaneously, we break the tie arbitrarily.) We continue this process until we reach a project which, when implemented, would overshoot the budget limit, and stop. A high-level difference to Rule X is that Rule X gives voters their share of the budget up front, while Phragmén's rule disburses shares over time. Phragmén's method fails EJR even for unit costs [Brill et al., 2017] and only satisfies the much weaker condition of PJR, so we focus on Rule X. In Appendix E we report on experiments in which we compare Rule X with Phragmén's rule and PAV when run on real data from several large participatory budgeting elections. Our experiments suggest that Rule X outperforms the other two rules with respect to several criteria pertaining to proportionality and efficiency.

[1] When formulated for unit costs, EJR can be applied recursively and thereby gains additional strength, since it implies that the number of well-represented voters in S is high [Sánchez-Fernández et al., 2017]: EJR requires that in S there exists a voter, say i, who is well-represented; if we remove this voter then the group $S \setminus \{i\}$ is still cohesive, and EJR would require that in this group there also exists a voter who is well-represented, etc.

FJR: A proportionality axiom even stronger than EJR

In approval-based multiwinner elections, it is fair to say that EJR is the strongest proportionality axiom that is known to always be satisfiable.[2] Many other rules such as Phragmén's rule or Chamberlin–Courant only satisfy substantially weaker axioms (known as PJR [Sánchez-Fernández et al., 2017] and JR [Aziz et al., 2017]).

A very attractive strengthening of EJR is the *core* [Aziz et al., 2017, Fain et al., 2018]. We say that a set $S \subseteq N$ of voters *blocks* an outcome W if there is a set T of projects affordable with a $|S|/|N|$ fraction of the budget such that each member of S strictly prefers T to W (in the sense that each member of S approves strictly more projects in T than in W). In such a case, the group S appears to be underrepresented. An outcome W is in the core if it is not blocked by any coalition S. It is unknown whether there always exists an outcome in the core in the approval-based model (even under the unit cost assumption), and this is surely the most tantalizing open problem in this area.

Note that the core implies EJR, since EJR requires that W is not blocked by a set T that is unanimous for S (i.e. all projects in T are approved by all members of S). We propose a property that is in between these two properties, by partially relaxing the cohesiveness requirement. We call this axiom Fully Justified Representation (FJR).[3] FJR requires that if a group $S \subseteq N$ of voters can propose a set T of projects that is affordable with S's share of the budget, and each voter approves at least ℓ projects in T, then at least one voter in S must approve at least ℓ projects in the chosen outcome W. Thus, rather than insisting that T is unanimously approved by the group S (like in EJR), we now allow cases where T is very popular among S, though not necessarily unanimous.

To the best of our knowledge, this natural axiom was not known to be satisfiable even for unit costs; in particular, both PAV and Rule X fail FJR (Examples 2 and 3). We prove that there does indeed exist a rule satisfying FJR, which works for arbitrary costs. The rule is further called Greedy Cohesive Rule (GCR)—it is a simple greedy procedure that repeatedly looks for groups with maximum cohesiveness and then satisfies them. While this is not a polynomial-time algorithm and not a particularly natural rule, we can show that this proposal can be made compatible with some other properties (in particular priceability, a property introduced by [Peters and Skowron, 2020]). For future work, it will be interesting to look for new natural rules satisfying FJR; this is even interesting for the committee context.

Beyond approval: Allowing more expressive preferences

In real-world PB elections, different projects differ vastly in their costs. For example, in the 2019 PB election in the 16th arrondissement of Paris, the most expensive project that was funded cost €560k (refurbishing a sports facility) and the cheapest cost €3k (providing materials for a school project of building a board game). The former project received 775 votes, and the latter 670 votes. The former project was 1.15 times as popular as the latter, but it cost 186 times as much! If we take the votes at face value, counting all approvals the same, it would seem that the cheap project provides an amazing value. It is more likely, though, that the approval-based interface did not allow voters to adequately express their values.

Facing these large cost differences, a better preference model might be given by general *additive valuations*, which allow voters to specify an arbitrary utility value for each project, with the assumption that a voter's satisfaction is proportional to the sum of the utilities of

[2] Though the literature contains other proportionality notions that are both logically and conceptually incomparable, such as "perfect representation" [Sánchez-Fernández et al., 2017] satisfied by the Monroe rule, and the concepts of laminar proportionality and priceability [Peters and Skowron, 2020] satisfied by Phragmén's rule and Rule X.

[3] Apologies that this name is not particularly descriptive, but then neither is EJR or PJR.

the funded projects. In the PB context, this model is considered by Benade et al. [2017] who study preference elicitation issues, and by Fain et al. [2018] and Fluschnik et al. [2019] who consider an aggregation rule similar to PAV, based on optimizing a Nash product objective. The latter rule will not satisfy us, given our discussion of Onetown and Twotown above. Further, even for unit costs, the rule does not satisfy our version of the EJR axiom. For Phragmén's rule, there seems to be no natural way at all to define it for general additive valuations. Also, the proportionality property PJR that Phragmén satisfies does not seem to have an analogue for general valuations.

For Rule X, however, we are able to propose a way to adapt it to general additive valuations. In our proposal, when Rule X decides to fund a project, a voter's payment is proportional to the voter's utility for the project. So if voter i assigns utility 1 to project B while j assigns utility $1/2$ to B, then Rule X will ask i to pay twice as much as j if B is funded. We also propose a natural way to extend the EJR axiom to general additive valuations. Rule X satisfies EJR *up to one project* – a mild relaxation. FJR can also be extended to general utilities, and our greedy rule satisfying this property continues to work for general additive utilities.

We close by discussing another input format, where voters have ordinal preferences, that is, where voters rank the projects in order of preference. We show that if we convert rankings into additive valuations using a lexicographic scheme, then our two voting rules give rise to voting rules for the ordinal setting. In particular, Rule X satisfies a property known as Proportionality for Solid Coalitions, a property first defined for the Single Transferable Vote (STV), a multi-winner voting rule used in many political constituencies in the anglosphere. GCR fails this property.

2 Preliminaries

For each $t \in \mathbb{N}$, write $[t] = \{1, 2, \dots, t\}$. An *election* is a tuple $(N, C, \mathrm{cost}, \{u_i\}_{i\in N})$, where:

- $N = [n]$ and $C = \{c_1, \dots, c_m\}$ $(n, m \in \mathbb{N})$ are the sets of *voters* and *candidates* (or *projects*), respectively.
- $\mathrm{cost}\colon C \to \mathbb{Q}_+$ is a function that for each candidate $c \in C$ assigns the *cost* that needs to be paid if c is selected. For each subset $T \subseteq C$, we write $\mathrm{cost}(T) = \sum_{c\in T} \mathrm{cost}(c)$ for the total cost of the projects in T.
- For each voter $i \in N$, the function $u_i\colon C \to [0, 1]$ defines i's additive *utility function*. If a set $T \subseteq C$ of candidates is implemented, i's overall utility is $u_i(T) = \sum_{c\in T} u_i(c)$. For a subset $S \subseteq N$ of voters, we further write $u_S(T) = \sum_{i\in S} \sum_{c\in T} u_i(c)$ for the total utility enjoyed by S if T is implemented. We assume that $u_N(c) > 0$ for each $c \in C$, that is, every candidate is assigned positive utility by at least one voter.

The voters have a fixed common *budget* which we normalize to 1. A subset of candidates $W \subseteq C$ is *feasible* if $\mathrm{cost}(W) \leqslant 1$. Our goal is to choose a feasible subset of candidates, which we call an *outcome*, based on voters' utilities. An *aggregation rule* (or, in short, a *rule*) is a function that for each election returns a set of feasible outcomes, called the *winning outcomes*.[4]

There are two interesting special cases of our model:

Committee elections. In this case, there exists $k \in \mathbb{N}$ such that each candidate costs $1/k$. Then W is an outcome if and only if $|W| \leqslant k$. In this special case we also refer to

[4]Typically, a rule will be selecting a single winning outcome, but ties are possible. For the results of this paper it won't matter how these ties are broken.

outcomes as *committees*, and we say that the election satisfies the *unit cost assumption*. For committee elections we will often refer to k as the maximal committee size.

Approval-based elections. In this case, for each $i \in N$ and $c \in C$ it holds that $u_i(c) \in \{0, 1\}$. The *approval set* of voter i is $A(i) := \{c \in C : u_i(c) = 1\}$, and we say that i *approves* candidate c if $c \in A(i)$. If $c \in A(i) \cap W$, we say that c is a *representative* of i.

Often we combine these special cases, and study approval-based committee elections.

3 Rule X

Recently, Peters and Skowron [2020] introduced an aggregation rule called *Rule X* for approval-based committee elections. In that setting, they showed that Rule X satisfies a combination of appealing proportionality properties. Here, we extend it to the more general model of participatory budgeting, that is, to the model with arbitrary costs and utilities.

Definition 1 (Rule X)**.** Each voter is initially given an equal fraction of the budget, i.e., each voter is given $1/n$ dollars. We start with an empty outcome $W = \emptyset$ and sequentially add candidates to W. To add a candidate c to W, we need the voters to pay for c. Write $p_i(c)$ for the amount that voter i pays for c; we will need that $\sum_{i \in N} p_i(c) = \mathrm{cost}(c)$. We write $p_i(W) = \sum_{c \in W} p_i(c) \leqslant \frac{1}{n}$ for the total amount i has paid so far. For $\rho \geqslant 0$, we say that a candidate $c \notin W$ is *ρ-affordable* if

$$\sum_{i \in N} \min\left(\tfrac{1}{n} - p_i(W), u_i(c) \cdot \rho\right) = \mathrm{cost}(c).$$

If no candidate is ρ-affordable for any ρ, Rule X terminates and returns W. Otherwise it selects a candidate $c \notin W$ that is ρ-affordable for a minimum ρ. Individual payments are given by

$$p_i(c) = \min\left(\tfrac{1}{n} - p_i(W), u_i(c) \cdot \rho\right) \qquad \square$$

Intuitively, when Rule X adds a candidate c, it asks voters to pay an amount proportional to their utility $u_i(c)$ for c; in particular, the cost per unit of utility is ρ. If a voter does not have enough money, the rule asks the voter to pay all the money the voter has left, which is $\frac{1}{n} - p_i(W)$. Throughout the execution of Rule X, the value ρ increases. Thus, candidates are added in decreasing order of utility per dollar that the voters get from the candidates.

Extended Justified Representation (EJR)

The first notion of proportionality that we examine is Extended Justified Representation (EJR). This axiom was first proposed for approval-based committee elections [Aziz et al., 2017]. Even for the special case of approval-based committee elections, only few rules are known to satisfy EJR [Aziz et al., 2017, 2018a, Peters and Skowron, 2020], but Rule X is one of them. In this section, we introduce a generalization of EJR to the PB model and show that our generalization of Rule X continues to satisfy EJR.

We first recall the definition of EJR for approval-based committee elections. Intuitively, this axiom ensures that every large enough group of voters whose approval sets have a large enough intersection must obtain a fair number of representatives.

Definition 2 (Extended Justified Representation for approval-based committee elections)**.** We say that a group of voters S is *ℓ-cohesive* for $\ell \in \mathbb{N}$ if $|S| \geqslant \ell/k \cdot n$ and $|\bigcap_{i \in S} A(i)| \geqslant \ell$.

A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E and each ℓ-cohesive group S of voters there exists a voter $i \in S$ such that $|A(i) \cap \mathcal{R}(E)| \geqslant \ell$.

At first sight it is unintuitive that we only require that at least one voter obtain ℓ representatives. However, the strengthening of EJR that requires each member of S to obtain ℓ representatives is impossible even on very small instances [Aziz et al., 2017]. Still, even with only the at-least-one guarantee, EJR has plenty of bite [Aziz et al., 2018a, Skowron, 2018, Peters and Skowron, 2020].

The generalization of this axiom to the PB model is not straightforward and to the best of our knowledge none has been proposed in the literature.[5] To warm up, let's first relax the unit cost assumption, but stay in the approval-based setting. Then EJR should state the following.

Definition 3 (Extended Justified Representation for approval-based elections)**.** We say that a group of voters S is T-cohesive for $T \subseteq C$ if $|S| \geqslant \mathrm{cost}(T) \cdot n$ and $T \subseteq \bigcap_{i \in S} A(i)$.

A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E and each T-cohesive group S of voters there exists a voter $i \in S$ such that $|A(i) \cap \mathcal{R}(E)| \geqslant |T|$.

Thus, cohesiveness now requires that the group S can identify a collection of projects T that they all approve and that is affordable with their fraction of the budget ($|S| \geqslant \mathrm{cost}(T) \cdot n$). Note that voters $i \in S$ obtain utility $u_i(T) = |T|$ from T; EJR requires that at least one member of S must attain this utility in the election outcome.

To further generalize EJR beyond approvals is more difficult, because the notion of a candidate who is approved by all members of S does not have an analogue. Instead, we quantify cohesion by calculating the minimum utility that any member of S assigns to each project in T.

Definition 4 (Extended Justified Representation)**.** A group of voters S is (α, T)-*cohesive*, where $\alpha \colon C \to [0;1]$ and $T \subseteq C$, if $|S| \geqslant \mathrm{cost}(T) \cdot n$ and if $u_i(c) \geqslant \alpha(c)$ for all $i \in S$ and $c \in T$.

A rule $\mathcal{R}$ satisfies *extended justified representation* if for each election instance E and each (α, T)-cohesive group of voters S there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geqslant \sum_{c \in T} \alpha(c)$.

Again, an (α, T)-cohesive group of voters S can propose the projects in T, since they are affordable with S's share of the budget. The values $(\alpha(c))_{c \in T}$ denote how much the coalition S agrees about the desirability of the projects in T. In particular, we have $u_i(T) \geqslant \sum_{c \in T} \alpha(c)$ for each $i \in S$. Consequently, Definition 4 prohibits any outcome in which every voter in S gets utility strictly lower than $\sum_{c \in T} \alpha(c)$; hence there must exists $i \in S$ such that $u_i(\mathcal{R}(E)) \geqslant \sum_{c \in T} \alpha(c)$.

We do not know if Rule X satisfies EJR in the general PB model. However, we can show that it satisfies a mild relaxation, which requires EJR to hold "up to one project".

Definition 5 (Extended Justified Representation Up To 1 Project)**.** A rule $\mathcal{R}$ satisfies extended justified representation *up to one project* if for each election instance E and each (α, T)-cohesive group of voters S there exists a voter $i \in S$ such that either $u_i(\mathcal{R}(E)) \geqslant \sum_{c \in T} \alpha(c)$ or for some $a \in C$ it holds that $u_i(\mathcal{R}(E) \cup \{a\}) > \sum_{c \in T} \alpha(c)$.

It is worth noting that in the approval-based model, Definitions 4 and 5 are actually equivalent, because the "up to one project" option never applies: Consider an (α, T)-cohesive group of voters S. Since voters' utilities are $0/1$, we may assume that for each $c \in T$ we have $\alpha(c) = 1$: if $\alpha(c) > 0$ this is clear; otherwise we can remove c from T without losing cohesiveness. Thus, the cohesiveness condition is equivalent to the condition that every voter approves every candidate in T. Finally, note that in the approval model, due to the strict

[5] Aziz et al. [2018b] generalize the weaker axiom of Proportional Justified Representation (PJR) [Sánchez-Fernández et al., 2017] beyond unit costs, but they operate in a non-standard utility model where voters care more about more expensive projects.

inequality, both conditions $u_i(\mathcal{R}(E)) \geqslant \sum_{c \in T} \alpha(c)$ and $\exists_{a \in C}. u_i(\mathcal{R}(E) \cup \{a\}) > \sum_{c \in T} \alpha(c)$ boil down to $|A(i) \cap \mathcal{R}(E)| \geqslant \sum_{c \in T} \alpha(c) = |T|$.

Our main result is that Rule X satisfies EJR up to one project, and hence it satisfies EJR in the approval-based model.

Theorem 2. *Rule X satisfies EJR up to one project in the participatory budgeting model.*

As we established in the introduction (Theorem 1), the most prominent rule that satisfies EJR for approval-based committee election does not extend this guarantee beyond unit costs. Let us briefly recall the definition of this rule.

Definition 6 (Proportional Approval Voting (PAV))**.** Consider an approval-based election. PAV selects a feasible outcome that maximizes $\sum_{i \in N} \mathrm{H}(|A(i) \cap W|)$, where $\mathrm{H}(r)$ is the r-th harmonic number, i.e., $\mathrm{H}(r) = \sum_{j=1}^{r} 1/j$.

Onetown, as shown in Figure 2, showed that PAV fails EJR (in the sense of Definition 3). Example 1 below constructs an alternative instance on which PAV does not satisfy EJR. In fact, this example shows that PAV does not even approximate EJR up to a constant factor.

Example 1. Fix a constant $r \in \mathbb{N}$ ($r \geqslant 2$), and consider the following approval-based profile:

$$\begin{array}{ll} r^2 - 1 \text{ voters:} & \{a_1, a_2, \ldots, a_r\}, \\ 1 \text{ voter:} & \{b_1, b_2, \ldots, b_r\}. \end{array}$$

The candidates $a_1, a_2, \ldots, a_r$ cost $1/r$ dollars each; the candidates $b_1, b_2, \ldots, b_r$ cost $1/r^3$ dollars each. EJR requires that the one voter who approves candidates $b_1, \ldots, b_r$ must approve at least r candidates in the outcome. However, PAV selects $\{a_1, a_2, \ldots, a_r\}$, leaving the voter with nothing. □

In Appendix B we discuss other properties of Rule X.

4 Greedy Cohesive Rule

In Section 3 we discussed the EJR axiom for the PB model, and saw that it is implemented by Rule X. One may wonder if there is natural strengthening of EJR that is still satisfiable. In this section we propose such a strengthening, and show that there is a rule that satisfies the new, stronger property. Interestingly, even in the approval-based committee-election model our new property is substantially stronger than EJR, and hence this new rule provides the strongest known proportionality guarantees. On the other hand, compared to Rule X, it is computationally more expensive and arguably less natural.

Full Justified Representation (FJR)

Our new proportionality axiom strengthens EJR by weakening its requirement that groups must be cohesive. Thus, the new axiom guarantees representation to groups that are only partially cohesive.

Definition 7 (Full Justified Representation (FJR))**.** We say that a group of voters S is *weakly* (β, T)*-cohesive* for $\beta \in \mathbb{R}$ and $T \subseteq C$, if $|S| \geqslant \mathrm{cost}(T) \cdot n$ and $u_i(T) \geqslant \beta$ for every voter $i \in S$.

A rule $\mathcal{R}$ satisfies *full justified representation (FJR)* if for each election instance E and each weakly (β, T)-cohesive group of voters S there exists a voter $i \in S$ such that $u_i(\mathcal{R}(E)) \geqslant \beta$.

In the approval-based committee-election model, FJR boils down to the following requirement: Let S be a group of voters, and suppose that each member of S approves at least β candidates from some set $T \subseteq C$ with $|T| \leqslant \ell$, and let $|S| \geqslant \ell/k \cdot n$. Then at least one voter from S must have at least β representatives in the committee. It is clear that in the special case of $\beta = \ell$, we obtain Definition 2, hence FJR implies EJR. The same implication holds in the general PB model.

Proposition 1. *FJR implies EJR in the general PB model.*

Proof. Suppose that rule $\mathcal{R}$ satisfies FJR and take an (α, T)-cohesive group of voters S for some $\alpha\colon T \to [0;1]$, $T \subseteq C$. For every voter $i \in S$ and every candidate $c \in T$ we have $u_i(c) \geqslant \alpha(c)$. We set $\beta = \sum_{c \in T} \alpha(c)$; clearly, we have also $u_i(T) \geqslant \beta$, thus S is weakly cohesive. As $\mathcal{R}$ satisfies FJR, we have that $u_i(\mathcal{R}(E)) \geqslant \beta = \sum_{c \in T} \alpha(c)$, which completes the proof. □

In turn, it is easy to see that FJR is implied by the core property (cf. Definition 9). It is related to, but stronger than, some other relaxations of the core discussed by Peters and Skowron [2020, Section 5.2]. The two major rules known to satisfy EJR for approval-based committee elections (Rule X and PAV) both fail FJR; to the best of our knowledge, no known rule satisfies FJR for approval-based committee elections, let alone for the general PB model.

Proposition 2. *Rule X and PAV do not satisfy FJR.*

Since no known aggregation rule satisfies FJR, one might wonder whether FJR existence can be guaranteed. It turns out that it can: we present a rule satisfying this strong notion of proportionality.

Definition 8 (Greedy Cohesive Rule (GCR))**.** The *Greedy Cohesive Rule* (GCR) is defined sequentially as follows: we start with an empty outcome $W = \emptyset$. At each step, we search for a value $\beta > 0$, a set of voters $S \subseteq N$, and a set of candidates $T \subseteq C$ such that S is weakly (β, T)-cohesive. If such values of β, S, and T do not exist, then we stop and return W. Otherwise, we pick values of β, S, and T that maximize β, breaking ties in favor of smaller $\mathrm{cost}(T)$.[6] We add all the candidates from T to W, remove all voters in S from the election and repeat the search.

Let us first check that the Greedy Cohesive Rule always selects a feasible outcome (i.e., that does not exceed the budget limit). Indeed, whenever the algorithm adds some set T to W, then by definition of weakly cohesive groups, we have $|S| \geqslant \mathrm{cost}(T) \cdot n$, and hence it removes at least $\mathrm{cost}(T) \cdot n$ voters after this step. Thus, if GCR selects an outcome with total cost $\mathrm{cost}(W)$, then it must have removed at least $\mathrm{cost}(W) \cdot n$ voters during its execution. Hence $\mathrm{cost}(W) \leqslant 1$.

Theorem 3. *Greedy Cohesive Rule satisfies FJR.*

Proof. Suppose that there exists a weakly (β, T)-cohesive group S which witnesses that FJR is not satisfied. Consider the voter $i \in S$ who was removed first by GCR and the outcome W right after that step (since S is weakly cohesive, such i always exists). Since $i \in S$ and S witnesses the FJR failure, we have $u_i(W) < \beta$. We know that i was removed as a member of some weakly (β', T')-cohesive group S'. Just before S' was removed, none of the members of S had been removed. Thus, we have $\beta' \geqslant \beta$, as GCR maximizes this value. However, since $T' \subseteq W$, we have $\beta' \leqslant u_i(T') \leqslant u_i(W) < \beta$—a contradiction to $\beta' \geqslant \beta$. Hence, such a group S does not exist. □

We discuss other properties of GCR in Appendix C.

[6] This way of breaking ties is important for our analysis of the priceability of GCR in Appendix C.1. The proof of Theorem 3 does not use it, so GCR satisfies FJR for any way of breaking ties.

5 Conclusion

In this paper, we have formulated two axioms, EJR and FJR, that capture the idea of proportionality in the participatory budgeting (PB) model. We have argued that none of the prominent committee election rules extend to the PB model so that it would satisfy even much weaker forms of proportionality. We have designed a simple and natural rule for the PB model, Rule X. It satisfies EJR and other proportionality-related properties such as priceability, and it is computable in polynomial time. The stronger of our two properties, FJR, is also satisfiable, albeit by a different and arguably less natural voting rule. It is an interesting open question whether there exists a natural voting rule that satisfies FJR and shares other desirable properties of Rule X.

In Appendix D we discuss how Rule X and GCR can be adapted for committee elections where voters have *ordinal preferences*, that is, voters express their preferences by ranking the candidates. We discuss properties of these two rules in the ordinal model.

In Appendix E, we present experiments comparing different variants of Rule X, Phragmén's rule, and PAV. These experiments are based on real data from participatory budgeting elections carried out in several major Polish cities (we looked at nine different elections). In our first experiment, we compared three different strategies of making Rule X exhaustive. We observed substantial differences between different variants of Rule X. We conclude that Rule X gives a lot of flexibility to a mechanism designer, because it often selects outcomes that do not spend all of the budget, while still satisfying strong fairness requirements like EJR. Depending on the specific objectives, a mechanism designer can choose to complete this outcome using different strategies. Among the strategies we described in Appendix B.3, we observed that the outcomes produced by Exh2 are better from a utilitarian perspective. It also divides the budget between different city districts in a substantially fairer way than outcomes produced by Exh1. Therefore we suggest Exh2 as the preferred method.

In our second experiment we compared Rule X, Phragmén's rule, and PAV. We observed that for approval-based utilities the results returned by Rule X and Phragmén's rule are comparably good, both in terms of the total utility obtained by the voters and in terms of proportionality. On the other hand, if we take a model with more fine-grained utilities, the difference between the two rules becomes apparent. This difference is unsurprising since Phragmén's rule does not take into account the more fine-grained information on utilities, but operates only on approval ballots. On the other hand, our results suggest that there is indeed a considerable advantage of using rules (like Rule X) that take into account the full information contained in cardinal additive utilities. We conclude that Rule X performs as well as Phragmén's rule for approval ballots and outperforms it when more fine-grained information on voters' utilities are available. Somewhat surprisingly, we show that the sequential variant of PAV produces highly disproportional outcomes compared to Phragmén's rule and Rule X. Altogether, our experiments confirm our theoretical results and suggest that Rule X outperforms the other two rules in terms of proportionality and/or efficiency.

As we discussed in the introduction, many cities run PB by dividing the overall budget between districts and running separate elections in each district. In particular, this is true in the Polish cities that provide our experimental election data. We claimed in the introduction that this practice of separate elections leads to inferior outcomes. We designed a final experiment to study this question. Our results show a visible advantage of using global rules such as Rule X over separate district elections. E.g., Rule X always produces outcomes with a more equal distribution of voter utility, and in most cases also provides a higher total utility in comparison to the rules that are in actual use in the elections we examined.

Acknowledgments Grzegorz Pierczyński and Piotr Skowron were supported by Poland's National Science Center grant UMO-2019/35/B/ST6/02215.

References

H. Aziz and B. Lee. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare*, 54(1):1–45, 2020.

H. Aziz and N. Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer, 2020.

H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sánchez-Fernández, and P. Skowron. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-2018)*, pages 902–909, 2018a.

H. Aziz, B. Lee, and N. Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*, pages 23–31, 2018b.

G. Benade, S. Nath, A. D. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*, pages 376–382, 2017.

M. Brill, R. Freeman, S. Janson, and M. Lackner. Phragmén's voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*, pages 406–413, 2017.

Y. Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.

Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. Group fairness in committee selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 263–279, 2019.

B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 575–592, 2018. Extended version arXiv:1805.03164.

P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, pages 27–47. AI Access, 2017.

P. Faliszewski, P. Skowron, S. Szufa, and N. Talmon. Proportional representation in elections: STV vs PAV. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2019)*, pages 1946–1948, 2019.

T. Fluschnik, P. Skowron, M. Triphaus, and K. Wilker. Fair knapsack. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-2019)*, pages 1941–1948, 2019.

S. Janson. Phragmén's and Thiele's election methods. Technical Report arXiv:1611.08826 [math.HO], arXiv.org, 2016. URL http://arxiv.org/abs/1611.08826.

Z. Jiang, K. Munagala, and K. Wang. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 463–472, 2020.

M. Lackner and P. Skowron. Consistent approval-based multi-winner rules. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC-2018)*, pages 47–48, 2018.

D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 2020 ACM Conference on Economics and Computation*, pages 793–794, 2020. Extended version arXiv:1911.11747.

D. Peters, G. Pierczynski, N. Shah, and P. Skowron. Market-based explanations of collective decisions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, 2021.

E. Phragmén. Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, 51(3):133–137, 1894.

E. Phragmén. *Proportionella val. En valteknisk studie.* Svenska spörsmål 25. Lars Hökersbergs förlag, Stockholm, 1895.

L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*, pages 670–676, 2017.

P. Skowron. Proportionality degree of multiwinner rules. Technical Report arXiv:1810.08799 [cs.GT], arXiv.org, 2018.

N. Talmon and P. Faliszewski. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-2019)*, pages 2181–2188, 2019.

T. N. Thiele. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pages 415–441. 1895.

N. Tideman and D. Richardson. Better voting methods through technology: The refinement-manageability trade-off in the Single Transferable Vote. *Public Choice*, 103(1–2):13–34, 2000.

D. Woodall. Properties of preferential election rules. *Voting Matters*, 3:Paper 4, 1994.

Dominik Peters
Harvard University
USA
Email: dpeters@seas.harvard.edu

Grzegorz Pierczyński
Department of Mathematics, Informatics and Mechanics
University of Warsaw
Poland
Email: g.pierczynski@mimuw.edu.pl

Piotr Skowron
Department of Mathematics, Informatics and Mechanics
University of Warsaw
Poland
Email: p.skowron@mimuw.edu.pl

A Omitted Proofs

A.1 Proof of Theorem 2

Theorem. *Rule X satisfies EJR up to one project in the participatory budgeting model.*

Proof. For the sake of contradiction, assume that there exists an election E, a subset of voters $S \subseteq N$ and a set of candidates $T \subseteq C$ such that: (i) $\mathrm{cost}(T) \leqslant |S|/n$, (ii) $u_i(c) \geqslant \alpha(c) > 0$ (candidates with $\alpha(c) = 0$ can be skipped) for each $i \in S$ and $c \in T$, and (iii) $u_i(\mathcal{R}(E) \cup \{a\}) \leqslant \sum_{c \in T} \alpha(c)$ for each $i \in S$ and $a \in T$.

Assume for a while that the voters from S have unrestricted initial budgets, and let us analyze how Rule X would proceed in such a case. For simplicity of presentation, without loss of generality, let us rename the candidates in T so that $T = \{c_1, \ldots, c_t\}$ and so that for $1 \leqslant i < j \leqslant t$ candidate c_i is picked by Rule X before candidate c_j.

Whenever a candidate $c \in T$ is selected, the voters pay for this candidate. Voter i pays $p_i(c)$ dollars for c, and in return, she gets utility $u_i(c)$. Thus, the price-per-utility she pays equals $\rho_i(c) = p_i(c)/u_i(c)$. Rule X works in a way that all voters from S who pay for c obtain the same price-per-utility ratio, i.e., for all $i, j \in S$ and $c \in C$ we have that $\rho_i(c) = \rho_j(c)$. Further, this price-per-utility equals at most $\mathrm{cost}(c)/u_S(c)$, independently of whether any voters from $N \setminus S$ pay for c or not (if no voters from $N \setminus S$ pays for c, then the price-per-utility equals exactly $\mathrm{cost}(c)/u_S(c)$):

$$\begin{aligned}\rho_i(c) &= \frac{p_i(c)}{u_i(c)} = \frac{p_i(c) \cdot \frac{\sum_{j \in S} u_j(c)}{\sum_{j \in S} u_j(c)}}{u_i(c)} = \frac{1}{\sum_{j \in S} u_j(c)} \cdot \sum_{j \in S} \frac{p_i(c)}{u_i(c)} \cdot u_j(c) \\ &= \frac{1}{\sum_{j \in S} u_j(c)} \cdot \sum_{j \in S} p_j(c) \leqslant \frac{\mathrm{cost}(c)}{u_S(c)}.\end{aligned} \tag{1}$$

Since $u_i(c) \geqslant \alpha(c)$ for each $i \in S$ and $c \in T$, the price-per-utility for $c \in T$ equals at most $\mathrm{cost}(c)/|S|\alpha(c)$. Now, consider the voter who in the first possible iteration uses more than $1/n$ dollars[7]. For this voter, call her i, let us consider the function f defined as follows. For each value x, the function f returns the price that i needs to pay to achieve the utility x. We make this function continuous, by assuming that the candidates are divisible. That is, if the voter pays p for her first paid candidate c with utility $u_i(c)$, then $f(u_i(c)/2) = p/2$, $f(u_i(c)/3) = p/3$, and so on. The key observation is that the function f is convex. This is because Rule X selects the candidates in increasing order of price-per-utility. This function is depicted below.

[7] The only case when there is no such voter is when every candidate $c \in C$ such that $u_S(c) > 0$ has already been elected before. In this case, the utility of i from the elected outcome is clearly at least $\sum_{c \in T} u_i(c) \geqslant \sum_{c \in T} \alpha(c)$.

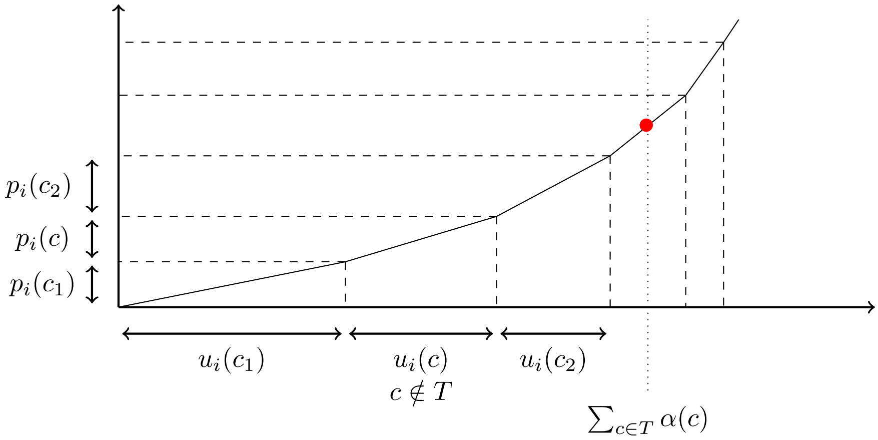

We are interested in the value $f(\sum_{c\in T} \alpha(c))$. This value would be maximized if the fragments of the function with the lowest slope were the shortest. However, we know that the part of the function that corresponds $\rho_i(c)$ must be of length at least equal to $u_i(c) \geqslant \alpha(c)$. Thus:

$$f\left(\sum_{c\in T} \alpha(c)\right) \leqslant \sum_{c\in T} \alpha(c) \cdot \rho_i(c) \leqslant \sum_{c\in T} \alpha(c) \cdot \frac{\mathrm{cost}(c)}{|S|\alpha(c)} = \sum_{c\in T} \frac{\mathrm{cost}(c)}{|S|} = \frac{\mathrm{cost}(T)}{|S|} \leqslant \frac{1}{n}.$$

Now, consider the first moment when i uses more than $1/n$ dollars. Until this time moment, Rule X behaves exactly in the same way as if the voters from S had their initial budgets set to $1/n$ (this follows from how we chose i). Further, we know that in this moment, if we chose a candidate that would be chosen if the voters had unrestricted budgets, then the utility of voter i would be greater than $\sum_{c\in T} \alpha(c)$. This gives a contradiction, and completes the proof. □

Proof of Proposition 2

Proposition. *Rule X and PAV do not satisfy FJR.*

The proof consists of two examples:

Example 2 (Rule X)**.** Consider the following instance of approval-based committee elections for $n = 22$ voters, $m = 13$ candidates, and where the goal is to select a committee of size $k = 11$:

voters 1-3:	$\{c_1, c_2, c_3, c_4, c_8\}$	voters 13-15:	$\{c_1, c_2, c_3, c_4, c_{12}\}$
voters 4-6:	$\{c_1, c_2, c_3, c_4, c_9\}$	voters 16-18:	$\{c_5, c_6, c_7, c_8, c_9, c_{10}, c_{11}, c_{12}\}$
voters 7-9:	$\{c_1, c_2, c_3, c_4, c_{10}\}$	voters 19-21:	$\{c_5, c_6, c_7\}$
voters 10-12:	$\{c_1, c_2, c_3, c_4, c_{11}\}$	voter 22:	$\{c_{13}\}$.

In the first 4 steps, Rule X chooses candidates c_1, c_2, c_3, c_4 (this happens for $\rho = 1/11{\cdot}15$). After that, each voter of the first 15 ones has $1/22 - 4/11{\cdot}15$ dollars. In next 3 steps, for $\rho = 1/11{\cdot}6$, candidates c_5, c_6, c_7 are chosen: 6 voters who support them spend all their money ($1/22 - 3/11{\cdot}6 = 0$). After that, the algorithm stops. Each of the first 15 voters has 4 candidates she approves; voters 16-18 approve 3 selected candidates. Thus, no member of the weakly $(5, \{c_1, c_2, c_3, c_4, c_8, c_9, c_{10}, c_{11}, c_{12}\})$-cohesive group of the first 18 voters has 5 representatives. □

Example 3 (PAV)**.** This example was first considered by Peters and Skowron [2020, Section 1]. We have $m = 15$ candidates and $n = 6$ voters, with the following preferences:

$$\begin{array}{llll} \text{voter } 1: & \{c_1, c_2, c_3, c_4\} & \text{voter } 4: & \{c_7, c_8, c_9\} \\ \text{voter } 2: & \{c_1, c_2, c_3, c_5\} & \text{voter } 5: & \{c_{10}, c_{11}, c_{12}\} \\ \text{voter } 3: & \{c_1, c_2, c_3, c_6\} & \text{voter } 6: & \{c_{13}, c_{14}, c_{15}\}. \end{array}$$

The size of the committee to be elected is $k = 12$. PAV chooses in this case committee $\{c_1, c_2, c_3, c_7, c_8, c_9, c_{10}, c_{11}, c_{12}, c_{13}, c_{14}, c_{15}\}$. Hence, no voter from the weakly $(4, \{c_1, c_2, c_3, c_4, c_5, c_6\})$-cohesive group consisting of the first 3 voters has 4 representatives. □

B Other Properties of Rule X

In this section we discuss other properties of Rule X, apart from EJR.

B.1 Approximating the Core

An important proportionality property advocated for PB is the core [Aziz et al., 2017, Fain et al., 2018]. It strengthens EJR by removing the cohesiveness requirement on groups S of voters.

Definition 9 (The Core)**.** For a given election instance $(N, C, \mathrm{cost}, \{u_i\}_{i \in N})$, we say that an outcome W is in the *core*, if for every $S \subseteq N$ and $T \subseteq C$ with $|S| \geqslant \mathrm{cost}(T) \cdot n$ there exists $i \in S$ such that $u_i(W) \geqslant u_i(T)$. We say that an election rule $\mathcal{R}$ satisfies the *core property* if for each election instance E the winning outcome $\mathcal{R}(E)$ is in the core.

While this is a clean and appealing notion, unfortunately, there are elections where no outcome is in the core, even with unit costs.

Example 4.[8] We have 6 voters and 6 candidates with unit costs, and $k = 3$. Utilities satisfy

$$\begin{array}{lll} u_1(c_1) > u_1(c_2) > 0, & u_2(c_2) > u_2(c_3) > 0, & u_3(c_3) > u_3(c_1) > 0; \\ u_4(c_4) > u_4(c_5) > 0, & u_5(c_5) > u_5(c_6) > 0, & u_6(c_6) > u_6(c_4) > 0, \end{array}$$

and all unspecified utilities are equal to 0. Let $W \subseteq C$ be any feasible outcome, so $|W| \leqslant 3$. Then either $|W \cap \{c_1, c_2, c_3\}| \leqslant 1$ or $|W \cap \{c_4, c_5, c_6\}| \leqslant 1$. Without loss of generality assume the former, and again without loss assume that $c_2 \notin W$ and $c_3 \notin W$. Then $S = \{v_2, v_3\}$ and $T = \{c_3\}$ block W and show it is not in the core, since $2 = |S| \geqslant \mathrm{cost}(T) * n = \frac{1}{3} \cdot 6 = 2$ and both $u_2(c_3) > u_2(c_1) \geqslant u_2(W)$ and $u_3(c_3) > u_3(c_1) \geqslant u_3(W)$. □

Notably, this example is not approval-based. It is unknown whether the core is always non-empty for approval-based elections (with or without the unit cost assumption).

In the committee context, Peters and Skowron [2020] showed that Rule X returns an outcome that never violates the core too badly; formally, Rule X provides a multiplicative approximation to the core.[9] We can generalize this result to the general PB setting: Rule X continues to provide a good multiplicative approximation to the core property.

Definition 10. For $\alpha \geqslant 1$, we say that an outcome is in the *α-core* if for every $S \subseteq N$ and $T \subseteq C$ with $\sum_{c \in T} \mathrm{cost}(c) \leqslant |S|/n$ there exists $i \in S$ and $c \in T$ such that $u_i(\mathcal{R}(E) \cup \{c\}) \geqslant \frac{u_i(T)}{\alpha}$.

[8] This example is adapted from Fain et al. [2018, Appendix C].

[9] Our approximation notion is different from the one proposed by Fain et al. [2018] and the one proposed by Cheng et al. [2019] and Jiang et al. [2020].

Theorem 4. *Given an election E, by $u_{\max}$ we denote the highest utility a voter can get from a feasible outcome. Analogously, by $u_{\min}$ we denote the smallest, yet still positive utility a voter can get from a feasible outcome:*

$$u_{\max} = \max_{i \in N} \max_{\mathrm{cost}(W) \leqslant 1} u_i(W) \qquad \text{and} \qquad u_{\min} = \min_{i \in N} \min_{u_i(W) > 0} u_i(W).$$

Rule X satisfies the α-core property for $\alpha = 4\log(2 \cdot {}^{u_{\max}}/_{u_{\min}})$.

Proof. Towards a contradiction, assume there exist an election instance E, a winning outcome $W \in \mathcal{R}(E)$, a subset of voters $S \subseteq N$, and a subset of candidates $T \subseteq C$ with $\sum_{c \in T} \mathrm{cost}(c) \leqslant {}^{|S|}/_{n}$ such that for each $i \in S$ and $c \in T$ it holds that $u_i(W \cup \{c\}) < {}^{u_i(T)}/_{\alpha}$.

Now, consider a fixed subset $S' \subseteq S$, and let:

$$\Delta(S') = \sum_{i \in S'} \big(u_i(T) - u_i(W)\big).$$

Similarly, as in the proof of Theorem 2, assume for a while that the voters from S' have unrestricted initial budgets, and let us analyze how Rule X would proceed in such a case. By the pigeonhole principle it follows that in each step of the rule there exists a not-elected candidate $c \in T \setminus W$ such that:

$$\frac{u_{S'}(c)}{\mathrm{cost}(c)} \geqslant \frac{\Delta(S')}{\mathrm{cost}(T)}.$$

Indeed, if for each $c \in T \setminus W$ we had $\frac{u_{S'}(c)}{\mathrm{cost}(c)} < \frac{\Delta(S')}{\mathrm{cost}(T)}$, then:

$$\Delta(S') \leqslant \sum_{c \in T \setminus W} u_{S'}(c) < \sum_{c \in T \setminus W} \mathrm{cost}(c) \cdot \frac{\Delta(S')}{\mathrm{cost}(T)} \leqslant \Delta(S'),$$

a contradiction.

Thus, the price-per-utility that the voters pay for the selected candidates equals at most $\frac{\mathrm{cost}(T)}{\Delta(S')}$. (This follows from the fact that Rule X selects candidates in such an order that the maximal price-per-utility the voters pay in a given round is minimized. The precise arguments are the same as in the proof of Theorem 2.) Now, consider the first moment when some voter in S'—call it i—uses more than its initial budget ${}^{1}/_{n}$. Until this time moment, Rule X bahaves exactly in the same way as if the voters from S' had their initial budgets set to ${}^{1}/_{n}$. Further, we know that in this moment, if we chose a candidate $c \in T$ that would be chosen if the voters had unrestricted budgets, then the voter i would pay more than ${}^{1}/_{n}$ in total, and thus, would get the utility of more than $\frac{1}{n} \cdot \frac{\Delta(S')}{\mathrm{cost}(T)}$. Since we assumed $u_i(W \cup \{c\}) < {}^{u_i(T)}/_{\alpha}$, we get that:

$$\frac{u_i(T)}{\alpha} > u_i(W) + u_i(c) > \frac{1}{n} \cdot \frac{\Delta(S')}{\mathrm{cost}(T)}.$$

Since $\alpha \geqslant 2$, and so $u_i(T) - u_i(W) \geqslant {}^{u_i(T)}/_{2}$, we get that:

$$u_i(T) - u_i(W) \geqslant \frac{u_i(T)}{2} > \frac{\alpha\Delta(S')}{2n \cdot \mathrm{cost}(T)}.$$

Let $S'' = S' \setminus \{i\}$. Clearly, we have that:

$$\Delta(S'') = \Delta(S') - (u_i(T) - u_i(W)) \leqslant \Delta(S')\left(1 - \frac{\alpha}{2n \cdot \mathrm{cost}(T)}\right).$$

The above reasoning holds for each $S' \subseteq S$. Thus, we start with $S' = S$ and apply it recursively, in each iteration decreasing the size of S' by 1. After $|S|/2$ iterations we are left with a subset S_e such that:

$$\Delta(S_e) \leqslant \Delta(S) \left(1 - \frac{\alpha}{2n \cdot \mathrm{cost}(T)}\right)^{\frac{|S|}{2}} \leqslant \Delta(S) \left(1 - \frac{\alpha}{2n \cdot \mathrm{cost}(T)}\right)^{\frac{\mathrm{cost}(T)n}{2}} < \Delta(S) \left(\frac{1}{e}\right)^{\frac{\alpha}{4}}.$$

Now, observe that $\Delta(S_e) \geqslant |S|/2 \cdot u_{\min}$ (for each $i \in S$ it must hold that $u_i(T) - u_i(W) > 0$) and that $\Delta(S) \leqslant |S| \cdot u_{\max}$. Thus, we get that:

$$\frac{|S|}{2} u_{\min} \cdot e^{\frac{\alpha}{4}} < |S| \cdot u_{\max},$$

which is equivalent to $e^{\frac{\alpha}{4}} < 2 \cdot \frac{u_{\max}}{u_{\min}}$ and, further, to $\alpha < 4 \log(2 \cdot u_{\max}/u_{\min})$. This gives a contradiction and completes the proof. □

The bound of $\alpha = 4 \log(2 \cdot u_{\max}/u_{\min})$ is asymptotically tight, and the hard instance can be constructed even for the approval-based committee-election model (there, $u_{\max}/u_{\min} \leqslant k$). The precise construction is given by Peters and Skowron [2020].

B.2 Priceability of Rule X

Peters and Skowron [2020] introduced a concept called priceability, which imposes a certain kind of balance on a voting rule. Every rule that, like Rule X, equally splits the budget between voters and then sequentially purchases projects using the money of its supporters will be priceable. Priceability does not place any restrictions on how the rule splits the project's cost among supporters. The concept also allows initial budgets higher than 1; an outcome is priceable if there exists *some* budget limit for which it is priceable.

Definition 11 (Priceability)**.** A price system is a pair $\mathsf{ps} = (b, \{p_i\}_{i \in N})$, where $b \geqslant 1$ is the initial budget (where each voter controls equal part of the budget, namely b/n), and for each voter $i \in N$, $p_i \colon C \to \mathbb{R}$ is a *payment function* that specifies the amount of money a particular voter pays for the elected candidates.[10,11] An outcome W is supported by a price system $\mathsf{ps} = (b, \{p_i\}_{i \in N})$ if the following conditions hold:

- **(C1).** $u_i(c) = 0 \implies p_i(c) = 0$ for each $i \in N$ and $c \in C$,[12]
- **(C2).** Each voter has the same initial budget of b/n dollars: $\sum_{c \in C} p_i(c) \leqslant b/n$ for each $i \in N$.
- **(C3).** Each elected candidate is fully paid: $\sum_{i \in N} p_i(c) = \mathrm{cost}(c)$ for each $c \in W$.
- **(C4).** The voters do not pay for non-elected candidates: $\sum_{i \in N} p_i(c) = 0$ for each $c \notin W$.
- **(C5).** For each unelected candidate $c \notin W$, the unspent budget of her supporters is at most $\mathrm{cost}(c)$: $\sum_{i \in N : u_i(c) > 0} \left(p - \sum_{c' \in W} p_i(c')\right) \leqslant \mathrm{cost}(c)$ for each $c \notin W$.

[10] Peters and Skowron [2020] assumed that each voter is initially given one dollar, which corresponds to setting $b = n$, but that there is an additional variable that specifies the total price that needs to be paid for an elected candidate. These two formulations are equivalent, but the present definition seems more natural for PB.

[11] The requirement that $b \geqslant 1$ ensures that the voters have at least enough money to buy candidates with total cost 1 (that is, the value of the real budget). Without this requirement, an empty outcome $W = \emptyset$ would be priceable (with $b = 0$).

[12] While condition (C1) is well-justified in the approval-based setting, in the general PB model it is very weak. Indeed, the condition does not put any restrictions on the payments when $u_i(c)$ is very small, yet positive.

An outcome W is said to be *priceable* if there exists a price system $\mathsf{ps} = (b, \{p_i\}_{i \in N})$ that supports W (that satisfies conditions (C1)–(C5)).

It is known that Rule X is priceable in the approval-based committee-election model and in the general PB model this property is also preserved—indeed, the rule implicitly constructs the price system satisfying the above conditions.

B.3 Exhaustiveness

A basic and very desirable efficiency notion is *exhaustiveness*, which requires that a voting rule spends its entire budget. Of course, due to the discrete model, we cannot guarantee that the rule will spend exactly 1 dollar (i.e., the entire budget); however, we can require that no additional project is affordable.

Definition 12 (Exhaustiveness, Aziz et al., 2018b)**.** An election rule $\mathcal{R}$ is *exhaustive* if for each election instance E and each non-selected candidate $c \notin \mathcal{R}(E)$ it holds that $\mathrm{cost}(\mathcal{R}(E) \cup \{c\}) > 1$.

Notably, Rule X fails to be exhaustive. It can happen that at the end of Rule X's execution, some project remains affordable, but the project's supporters do not have enough money to pay for it − Rule X then refuses to fund the project. For example, if we have two voters and two candidates, such that v_1 approves $\{c_1\}$ and v_2 approves $\{c_2\}$. Suppose that both candidates cost 1 dollar. Then Rule X returns $W = \emptyset$. In fact, it turns out that exhaustiveness is incompatible with priceability.

Example 5. We have 3 candidates and 3 voters. The first 2 voters approve $\{c_1\}$, and the third one approves $\{c_2, c_3\}$. We have $\mathrm{cost}(c_1) = 1$ and $\mathrm{cost}(c_2) = \mathrm{cost}(c_3) = 1/3$. The only exhaustive outcomes are $\{c_1\}$ and $\{c_2, c_3\}$. However, neither of them is priceable—indeed, to buy both c_2 and c_3, the third voter needs to control at least $2/3$ dollars. Then the first two voters control $4/3$ dollars and can buy candidate c_1, a contradiction. On the other hand, to buy c_1, the first two voters need to control at least 1 dollar. Then, the third voter controls at least $1/2$ dollars and buys c_2 or c_3, a contradiction. □

In some contexts, it may actually be a desirable feature of Rule X that it is not exhaustive, especially if unspent budget can be used in other productive ways (such as in next year's PB election). Arguably, in non-exhaustive examples, no remaining project has sufficient support to justify its expense; on that view, no further projects should be funded. In other situations, unspent budget may not be reusable, such as when the budget comes from a grant where unspent money needs to be returned (and the relevant decision makers do not obtain value from the grant-maker's alternative activities), or when the 'budget' is time (for example, when we use PB to plan activities for a day-long company retreat). In such situation, one might prefer an exhaustive rule.

Peters and Skowron [2020] proposed to complete the outcome elected by Rule X with the use of Phragmén's sequential rule (with initial budgets of the voters equal to the remainder left after running Rule X). We defined Phragmén's rule in the introduction. However, there is no obvious way of generalizing Phragmén's rule to non-approval utilities.[13]

Since we have generalized Rule X to work for general additive valuations, there is another way for us to make it exhaustive. Recall that Rule X fails to be exhaustive in situations

[13] One possibility, similar to our generalization of Rule X, would be to require that voters' payment for selected projects must be proportional to their utilities. Interestingly, this idea seems to not work at all for Phragmén's rule. For example, consider a committee election with two projects, c_1 and c_2, committee size $k = 1$, and two voters: The first voter assigns utility 1 to c_1 and utility 100 to c_2; the second voter assigns utility 1 to c_1 and 99 to c_2. When forced to use proportional payments, Phragmén's rule would choose c_1, a very inefficient choice.

where the remaining projects' supporters do not have sufficient funds left. However, in elections where $u_i(c) > 0$ for all $i \in N$ and $C \in C$, every voter supports every candidate, and thus this problem never occurs. In fact, Rule X is exhaustive when run on profiles of this type.

Proposition 3. *Consider an election $E = (N, C, \mathrm{cost}, \{u_i\}_{i \in N})$ such that $u_i(c) > 0$ for each $i \in N$ and $c \in C$. The outcome returned by Rule X for E is exhaustive.*

Proof. For the sake of contradiction assume that an outcome W returned by Rule X for an election instance E is not exhaustive. Then, there exists a candidate $c \notin W$ such that $\mathrm{cost}(W \cup \{c\}) \leqslant 1$. The voters paid in total $\mathrm{cost}(W)$ dollars for W; their initial budget was 1, thus after W is selected they all have at least $\mathrm{cost}(c)$ unspent money. However, this means that at the end of the execution of Rule X there exists a possibly very large value of ρ such that:

$$\sum_{i \in N} \min\left(\tfrac{1}{n} - p_i(W), u_i(c)\rho\right) = \sum_{i \in N} \left(\tfrac{1}{n} - p_i(W)\right) \geqslant \mathrm{cost}(c).$$

Consequently, c (or some other candidate) would be selected by Rule X, a contradiction. □

Thus, we can make Rule X exhaustive by perturbing the input utilities so that all utility values are positive; we call this strategy Exh1. Specifically, for a small $\varepsilon > 0$ ($\varepsilon \ll \min_{u_i(c)>0} u_i(c)$), and for each $i \in N, c \in C$ such that $u_i(c) = 0$ in the initial instance we set $u_i^{\varepsilon}(c) = \varepsilon$. Next, we run Rule X on the modified instance $\{u_i^{\varepsilon}\}_{i \in N}$; by Proposition 3 the outcome is exhaustive. Finally we return the outcome identified by Rule X as $\varepsilon \to 0$; the result is well-defined by the following result.

Proposition 4. *Consider any election $E = (N, C, \mathrm{cost}, \{u_i\}_{i \in N})$. There exists some $\overline{\varepsilon} > 0$ such that for all $0 < \varepsilon_1, \varepsilon_2 < \overline{\varepsilon}$, Rule X returns the same outcome when run on $\{u_i^{\varepsilon_1}\}_{i \in N}$ and $\{u_i^{\varepsilon_2}\}_{i \in N}$.*

This process gives rise to a different voting rule, an exhaustive variant of Rule X. Note that this rule is not priceable since it may ask voters to pay for candidates that they assign utility 0. By necessity, this rule sometimes elects candidates that cannot be afforded by its supporters. In these cases, when we elect such a candidate c, we will ask all supporters of c to pay all their remaining money for c, and split the remaining cost to be paid equally among voters who do not support c. Say that the maximum amount paid by a non-supporter for c is x; Rule X selects those candidate that minimize the value x at each step. Because voters are asked to spend their entire remaining budget if a non-affordable candidate they like is elected, this extension of Rule X will not distort the outcome too much.

Peters et al. [2021] suggested yet another method of completing outcomes returned by Rule X. This method, which we call Exh2, was proposed in the context of approval ballots, but it naturally extends to cardinal utilities. In Exh2 we run Rule X with the initial value of the budget b_i set to a value possibly greater than the actual available budget, i.e., $b_i \geqslant 1$. Using a binary search we find the highest value of b_i such that the total cost of the projects selected by Rule X does not exceed the actual budget. This method does not guarantee that the budget is exhausted (for a detailed discussion, see [Peters et al., 2021]), but in most cases this is indeed the case.

In Appendix E.1 we experimentally compare Exh1 and Exh2 based on data collected from real participatory budgeting instances. Our experiments show that Exh2 typically returns outcomes with higher values of total voters' satisfaction, and we suggest it as the preferred method.

C Additional Results about the Greedy Cohesive Rule

C.1 Priceability and Exhaustiveness of GCR

GCR satisfies neither priceability (Definition 11) nor exhaustiveness (Definition 12). However, in Appendix C.1, using a generalization of Hall's marriage theorem, we prove that the outcome of GCR can always be completed to a priceable outcome. We also show that GCR can be made exhaustive in a way that is almost priceable. In Appendix C.2, we discuss some drawbacks of GCR by discussing two examples where GCR arguably selects a bad outcome.

We start by proving three useful lemmas.

Lemma 1 (Polyandrious generalization of Hall's marriage theorem)**.** *Let $G = (U + V, E)$ be a bipartite graph and for every $A \subseteq U$ denote by $N_G(A)$ the neighbourhood of A, i.e. $N_G(A) = \{v \in V : \exists u \in A. \{u, v\} \in E\}$. Let $q \in \mathbb{N}$. Then for each $A \subseteq U$ we have that $|N_G(A)| \geqslant |A| \cdot q$ if and only if there exists a one-to-q mapping from each vertex in U to some q vertices in V, such that to each vertex from V at most one vertex from U is mapped.*

Proof. Consider the graph G' obtained by replacing set U with its q copies: $U_1, \ldots, U_q$ (we also copy edges). Consider now any set $A \in \bigcup_i U_i$. As $\bigcup_i U_i$ consists of q separate copies, there exists $i \in [q]$ such that $|A \cap U_i| \geqslant {}^{|A|}/_{q}$. Hence, from our assumption we have that $|N_G(A)| \geqslant |N_G(A \cap U_i)| \geqslant q|A \cap U_i| \geqslant |A|$. Now, from Hall's marriage theorem, in G' there exists a matching between $\bigcup_i U_i$ and V, covering set $\bigcup_i U_i$. Hence, in graph G it is enough to map each vertex $u \in U$ to these q vertices in V, to which q copies of u are matched in G'. Naturally, the implication holds also in the reverse direction—if there exists a one-to-q mapping in G as described above, then trivially for all $A \subseteq U$ we have that $|N_G(A)| \geqslant q|A|$. ☐

Lemma 2. *Let S be n (β, T)-cohesive group which is selected in some step of GCR. For every subset $A \subseteq T$, the size of the set of voters $S' := \{i \in S : u_i(A) > 0\}$ is at least $\mathrm{cost}(A) \cdot n$.*

Proof. The statement is trivial for $\mathrm{cost}(A) = 0$, so assume that $\mathrm{cost}(A) > 0$. Suppose for the sake of contradiction that the set $S' \subseteq S$ defined above has smaller size than $\mathrm{cost}(A) \cdot n$. Then group $S \setminus S'$ together with set $T \setminus A$ is $(\beta, T \setminus A)$-cohesive. Indeed,

$$|S \setminus S'| \geqslant |S| - \mathrm{cost}(A) \cdot n \geqslant \mathrm{cost}(T) \cdot n - \mathrm{cost}(A) \cdot n = \mathrm{cost}(T \setminus A) \cdot n.$$

Thus, as $\mathrm{cost}(A) > 0$, we have $\mathrm{cost}(T \setminus A) < \mathrm{cost}(T)$. Thus, GCR would select $S \setminus S'$ instead of group S, a contradiction. ☐

Lemma 3. *For every outcome of the GCR rule, there always exists a payment function satisfying conditions (C1)–(C4) with $b = 1$.*

Proof. Consider a single step of GCR and let S be an (β, T)-cohesive group considered in that step. We will prove that there exists a price system in which voters from S pay $\mathrm{cost}(c)$ dollars for each candidate $c \in T$[14].

Denote by d the least common multiple of the denominators of the rational numbers from the set: $\{\mathrm{cost}(c) : c \in T\}$. Note that ${}^{1}/_{d}$ is a divisor of all these costs. Assume that each candidate c is splitted into $\mathrm{cost}(c) \cdot d$ parts, each one associated with cost ${}^{1}/_{d}$. Besides, assume that each voter has d *coins*, each one worth ${}^{1}/_{d \cdot n}$ dollars.

[14]Note that if a candidate $c \in T$ has been considered in previous steps of the algorithm, she does not need to be paid again. However, this case would even strengthen the proof (we could just not charge voters assigned to paying for her and conditions (C1)–(C4) would be satisfied), so further we assume that T contains only candidates not elected yet.

Consider the bipartite graph $G = (A_T + A_S, E)$, where A_S is the set of all voters' coins and A_T is the set of all candidates' parts. In G there is an edge between a coin of a voter $i \in S$ and a part of a candidate $c \in T$ if and only if $u_i(c) > 0$.

Now, consider a set $A \subseteq A_T$, and let us assess the size of the neighbourhood $N_G(A)$. Let $C(A)$ denote the set of candidates whose some parts belongs to A. There is an edge from $p \in A$ to a coin of a voter i only if i assigns a positive utility to the candidate of p. Thus, $N_G(A)$ consists of coins of those voters, who assign a positive utility to some candidate from $C(A)$. By Lemma 2 there are at least $\mathrm{cost}(C(A)) \cdot n$ such voters, each voter comes with d coins, thus:

$$|N_G(A)| \geqslant \mathrm{cost}(C(A)) \cdot n \cdot d \geqslant \mathrm{cost}(A) \cdot n \cdot d.$$

Further, since each part of A costs exactly $1/d$, we get that:

$$|N_G(A)| \geqslant \mathrm{cost}(A) \cdot n \cdot d = 1/d \cdot |A| \cdot n \cdot d = |A| \cdot n.$$

Hence, from Lemma 1 we have that there exists a mapping from A_T to A_S such that every part of every candidate c is mapped to n coins and to each coin at most one candidate is mapped.

Now the payment function is constructed as follows: for every voter $i \in S$ and candidate $c \in T$, if exactly q coins of i are mapped with some parts of c, then $p_i(c) = q/d \cdot n$. It is straightforward to check that such a payment function satisfies conditions (C1)–(C4) for $b = 1$, which completes the proof. □

Finally, we can state the main result of this subsection.

Theorem 5. *Every outcome W elected by GCR can be completed to some priceable outcome.*

Proof. From Lemma 3, we know that there exists a family of payment functions $\{p\}_{i \in N}$ satisfying conditions (C1)–(C4) for W. Now, to obtain outcome W' supported by a valid price system, it is enough to run Rule X for this instance with initial outcome set to W and the initial endowment of every voter $i \in N$ set to $1/n - \sum_{c \in W} p_i(c)$. □

C.2 Some Drawbacks of the Greedy Cohesive Rule

Since GCR satisfies FJR but Rule X does not, we may conclude that GCR is a better rule. Clearly, GCR is custom-engineered to satisfy FJR. Thus, we may expect the rule to be deficient in other dimensions. The results presented in Appendix C.1 certainly suggest that GCR is not pathological, but in this section we consider some examples where Rule X seems to select better outcomes than GCR.

We begin by discussing a property that Peters and Skowron [2020] call *laminar proportionality*. This property identifies a family of well-behaved preference profiles and specifies the outcome on those profiles. The axiom is defined for the case of approval-based committee elections. Rule X satisfies; the following example shows that GCR does not.

Example 6 (GCR fails laminar proportionality)**.** Let $N = \{1, 2, 3, 4\}$ and $k = 8$, and introduce the candidate sets $X = \{c_1, \dots, c_4\}$, $Y = \{c_5, \dots, c_{10}\}$, and $Z = \{c_{11}, c_{12}\}$. The first three voters approve $X \cup Y$, and the fourth one approves $X \cup Z$. Two copies of the profile are depicted below. The candidates are represented by boxes; each candidate is approved by the voters who are below the corresponding box.

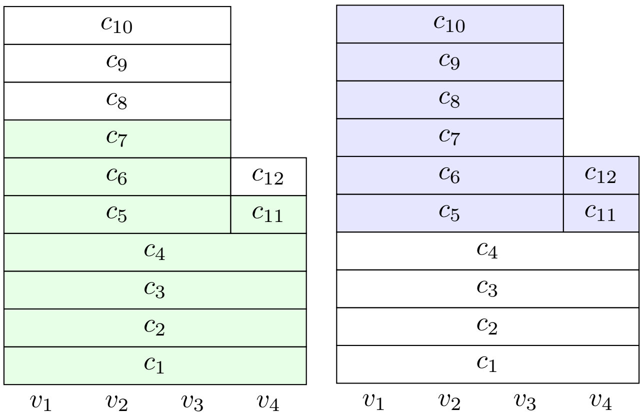

In this election instance, laminar proportionality would require that the voting rule selects all the candidates from X since they are approved by everyone. After electing the candidates in X, four seats are left to fill. Since the group $\{v_1, v_2, v_3\}$ the three times as large as the group $\{v_4\}$, laminar proportionality requires that we elect three candidates from Y and one candidate from Z. Thus, the committee indicated by the green boxes on the left-hand figure is laminar proportional.

On the other hand, in the first step GCR can choose the weakly $(6, Y)$-cohesive group $\{v_1, v_2, v_3\}$ and in the second step it can select the weakly $(2, Z)$-cohesive group $\{v_4\}$. This results in the blue committee depicted in the right-hand figure; this committee fails laminar proportionality. $\square$

Example 6 shows that in general, GCR is not laminar proportional, as it can return committees which are prohibited by the axiom. However, this example is not fully satisfactory, as it depends on tie-breaking. For example, in the first step we could choose the weakly $(6, \{c_2, c_3, c_4, c_5, c_6\})$-cohesive group containing the first three voters, and in the second step the weakly $(2, \{c_1, c_{11}\})$-cohesive group containing the last voter. An open question is whether GCR can always elect a committee satisfying laminar proportionality (among others). However, the following example shows that for some 'nearly laminar' instances, GCR does not match the general intuition standing behind this axiom.

Example 7. Modify the instance described in Example 6 in the following way: we have $N = [4000]$. Voter 1 approves only candidates from Y, voters 2 to 3000 approve $X \cup Y$, voters 3001 to 3999 approve $X \cup Z$ and voter 4000 approves Z.

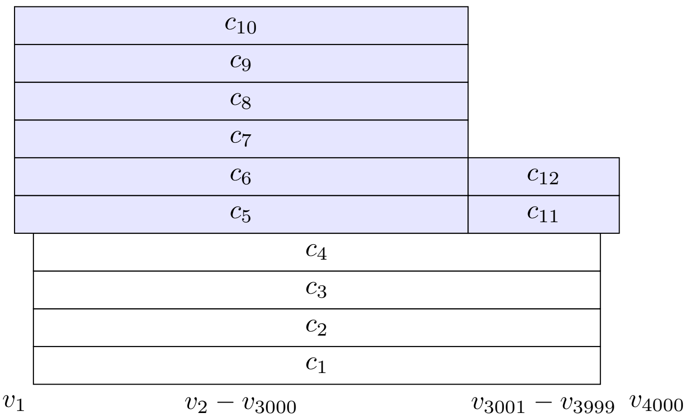

This instance is not laminar (because of the two voters not approving X), but it is close to being laminar and it is reasonable to expect that the elected committee should be the same as the green one from Example 6. Rule X uniquely elects that committee. On the other hand, GCR selects first the weakly $(6, Y)$-cohesive group containing the first 3000 voters and in the second step the weakly $(2, Z)$-cohesive group containing the last 1000 voters. After that the algorithm stops, electing committee $Y \cup Z$, as depicted above. Note that in this case, the choice of weakly cohesive groups is unique. $\square$

Examples 6 and 7 do not rule out the existence of an FJR rule that is also laminar proportional; the existence of a natural example of such a rule is an interesting open problem.

D Rule X and GCR for Ordinal Ballots

In this section we discuss how our two rules can be adapted for committee elections where voters have *ordinal preferences*, that is, voters express their preferences by ranking the candidates. The main idea is straightforward: we convert voters' preference rankings into additive valuations, by using positional scoring rules, and then apply our rules to the resulting election. Note that if we use scoring rules that assign positive values to positions in voters' rankings, we always obtain exhaustive rules. We will show that when we use a lexicographic conversion scheme in which voters care infinitely more about their top-ranked candidate than their second-ranked candidate and so on, then Rule X satisfies an axiom called Proportionality for Solid Coalitions (PSC). which was first introduced to analyze the Single Transferable Vote [Woodall, 1994, Tideman and Richardson, 2000]. (GCR does not satisfy PSC.) Rule X as applied to ordinal preferences is related to the Expanding Approvals rule of Aziz and Lee [2020]. Interestingly, due to its flexibility, Rule X can be used to extend the proportionality idea behind PSC beyond a lexicographic interpretation of preferences: Depending on how we convert voters' preference rankings to cardinal utilities, we obtain different forms of proportionality (cf., Faliszewski et al., 2019). For example, if we use Borda scores, the rule chooses outcomes where the average position of selected candidates in voters' rankings is high.

D.1 Model for Ordinal Preferences

In this section we assume that each voter $i \in N$ submits a strict preference order $\succ_i$ over the set of candidates. The order $c_{i_1} \succ_i c_{i_2} \succ \ldots \succ c_{i_m}$ means that c_{i_1} is voter's i most preferred candidate, c_{i_2} is her second most-preferred candidate, and so on. By $\mathrm{pos}_i(c)$ we denote the position of candidate c in i's preference ranking. In the above example we have $\mathrm{pos}_i(c_{i_1}) = 1$, $\mathrm{pos}_i(c_{i_2}) = 2$, and so on. For sets A and B, we write $A \succ_i B$ if $a \succ_i b$ for all $a \in A, b' \in B$.

Further, we assume unit costs, so that the goal is to select a committee of k candidates, and thus that the cost of each candidate is $1/k$.

Definition 13 (Proportionality for Solid Coalitions (PSC))**.** An outcome W satisfies PSC if for each $\ell \in [k]$, each subset of voters $S \subseteq N$ with $|S| \geqslant n\ell/k$, and each subset of candidates T such that $T \succ_i C \setminus T$ for all $i \in S$, it holds that $|W \cap T| \geqslant \min(\ell, |T|)$.

A rule satisfies PSC if for each election it only returns outcomes that satisfy PSC.

Definition 13 focuses on voters' top preferences—intuitively, it requires that if $c \succ_i c'$, then the utility that voter i assigns to candidate c is infinitely higher than that assigned to c'. Rule X naturally extends to such preferences, which we call *lexicographic utilities*, but we need to adapt Definition 1 to use a slightly different interpretation of the price per unit of utility, ρ. So far we assumed that ρ is a positive real value; in order to adapt the definition to lexicographic preferences we assume that $\rho \in [m]$, and that the multiplication by candidates' utilities is defined as follows:

$$\rho \cdot u_i(c) = \begin{cases} 1 & \text{if } \rho \geqslant \mathrm{pos}_i(c), \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 5. *Rule X for lexicographic utilities satisfies PSC.*

Proof. Consider a committee W returned by Rule X for an election instance $(N, C, k, \{\succ_i\}_{i \in N})$. Let $\ell \in [k]$, $S \subseteq N$, and T be as in Definition 13. The voters in S initially have the following budget:

$$|S| \cdot {}^{1}\!/_{n} \geqslant {}^{n\ell}\!/_{k} \cdot {}^{1}\!/_{n} = {}^{\ell}\!/_{k}.$$

Consider the steps of Rule X as the price per unit of utility, ρ, increases from 1 to $|T|$. In each such step, each voter from S can pay only for the candidates in T. Indeed, each candidate $c \in C \setminus T$ occupies a worse position than $|T|$ in those voters' preference rankings, and so for each $i \in S$ we have $\rho \cdot u_i(c) = 0$ (since $\rho \leqslant |T|$). When ρ reaches $|T|$, then for each candidate $c \in T$ and each $i \in S$ we have $u_i(c) = 1$. The voters from S have enough money to buy ℓ candidates, and so they will buy at least $\min(\ell, |T|)$ candidates from T. □

One may wonder, given the lexicographic utility scheme, whether PSC is just a consequence of Rule X satisfying EJR. Example 8 below shows that this is not the case and that the two axioms are logically incomparable in this context. FJR and PSC are also logically incomparable.

Example 8. Consider three voters with the following preference orders over $C = \{c_1, c_2, c_3, c_4\}$:

$$\begin{aligned} 1&\colon c_1 \succ c_2 \succ c_3 \succ c_4 \\ 2&\colon c_2 \succ c_3 \succ c_1 \succ c_4 \\ 3&\colon c_3 \succ c_1 \succ c_2 \succ c_4. \end{aligned}$$

Assume $k = 2$. In this example PSC would require that two candidates from $\{c_1, c_2, c_3\}$ are elected. On the other hand, committee $\{c_1, c_4\}$ satisfies FJR.

Now, consider two voters with the following preferences:

$$\begin{aligned} 1&\colon c_1 \succ c_2 \succ c_3 \succ c_4 \\ 2&\colon c_4 \succ c_1 \succ c_3 \succ c_2. \end{aligned}$$

Assume $k = 1$. Here, EJR requires that c_1, c_2, or c_4 must be selected. On the other hand, $\{c_3\}$ is a committee that satisfies PSC. □

GCR can also be adapted to lexicographic utilities. In this case, it is sufficient to assume that the utilities are exponentially decreasing with the positions—for each $i \in N$ and $c \in C$ we set $u_i(c) = m^{-\mathrm{pos}_i(c)}$. Then, for each c we have that $u_i(c) > \sum_{c' \prec_i c} u_i(c')$, and so the utility a voter assigns to a candidate in position p is higher than the utility that it would assign to any committee all of whose members are ranked below p.

Proposition 6. *GCR for lexicographic utilities fails PSC.*

Proof. We show that GCR fails PSC. Consider the following preference profile:

$$
\begin{aligned}
1 &: c_1 \succ c_7 \succ c_8 \succ c_6 \succ c_4 \succ c_5 \succ c_2 \succ c_3 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
2 &: c_1 \succ c_7 \succ c_8 \succ c_6 \succ c_4 \succ c_5 \succ c_2 \succ c_3 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
3 &: c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_4 \succ c_5 \succ c_7 \succ c_8 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
4 &: c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_4 \succ c_5 \succ c_7 \succ c_8 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
5 &: c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_4 \succ c_5 \succ c_7 \succ c_8 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
6 &: c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_4 \succ c_5 \succ c_7 \succ c_8 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
7 &: c_2 \succ c_3 \succ c_1 \succ c_7 \succ c_8 \succ c_4 \succ c_5 \succ c_6 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
8 &: c_3 \succ c_2 \succ c_1 \succ c_7 \succ c_8 \succ c_4 \succ c_5 \succ c_6 \succ c_9 \succ c_{10} \succ c_{11} \succ c_{12}\\
9 &: c_4 \succ c_5 \succ c_9 \succ c_7 \succ c_8 \succ c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_{10} \succ c_{11} \succ c_{12}\\
10 &: c_5 \succ c_4 \succ c_9 \succ c_7 \succ c_8 \succ c_1 \succ c_2 \succ c_3 \succ c_6 \succ c_{10} \succ c_{11} \succ c_{12}\\
11 &: c_{10} \succ c_{11} \succ c_{12} \succ c_7 \succ c_8 \succ c_1 \succ c_2 \succ c_3 \succ c_4 \succ c_5 \succ c_6 \succ c_9\\
12 &: c_{11} \succ c_{10} \succ c_{12} \succ c_7 \succ c_8 \succ c_1 \succ c_2 \succ c_3 \succ c_4 \succ c_5 \succ c_6 \succ c_9.
\end{aligned}
$$

Assume $k = 4$. Here, GCR will first pick $S = \{1, \dots, 6\}$ as a weakly cohesive group, with the corresponding set of candidates $T = \{c_1, c_6\}$. Indeed, if T consisted of 3 candidates, then S would need to have at least 9 voters. However, any 9 voters rank at least 4 different candidates at the top position, thus at least one of them would have a lower satisfaction than the voters from S have from T. By the same argument, T cannot consist of 4 candidates. If T consisted of 2 candidates but S included one voter from $7, \dots, 12$, then the satisfaction of voter 1 or 2 would also be lower. Indeed, these two voters rank c_2, c_3, c_4, c_5, c_{10}, and c_{11} (that is candidates that appear in the top positions) below c_6.

Hence, GCR picks c_1 and c_6, and removes the first 6 voters from further consideration. In the second step, the rule picks c_7 and c_8. This is because each other candidate appears at most twice before c_7 and c_8 in the remaining voters' rankings. Thus, the rule picks c_1, c_6, c_7 and c_8.

On the other hand, by looking at voters $3, \dots, 8$ we observe that PSC requires that two candidates from c_1, c_2, c_3 should be selected. □

E Experiments

In this section we look at data describing voters' preferences, collected from participatory budgeting elections carried out in several major cities in Poland. Based on this data we make a number of observations regarding election rules that we discuss in this paper.

In election instances that we have considered the projects were divided into several groups. One group consisted of city-wide projects, and each other group contained projects that were assigned to one of several city districts. Each voter was allowed to approve at most ten city-wide projects, and at most ten projects from her district. A part of the municipal budget was assigned to city-wide projects and the other part was divided among the districts, proportionally to their populations. Currently, the cities that we consider use a rule that selects projects greedily until the budget is exhausted; the projects are picked in the order of the number of garnered approval votes.

In our experimental analysis we used two types of voters' preferences:

Approval utilities: corresponding directly to the approval-ballots from our PB data.

Cardinal utilities: for each voter i and each project c_j we obtained the value of the cardinal utility $u_i(c_j)$ as follows. If i does not approve c_j, we assume that $u_i(c_j) = 0$.

If i approves c_j, we sample $u_i(c_j)$ from the normal distribution centred at $\mathrm{cost}(c_j)$. Intuitively, this means that voters care more about (take a greater utility from) funding expensive projects they like. This assumption is consistent with properties considered by Talmon and Faliszewski [2019], and with the metod of preference elicitation called "value-for-money" [Benade et al., 2017]. We also tested similar assumptions, where the values $u_i(c_j)$ were sampled from the uniform and exponential distributions, but the results led to qualitatively consistent conclusions. We also checked the case, where the utilities were sampled from distributions centred at a constant value instead of values dependent on the costs of the projects; in such cases, the results were similar to those that we observed for approval-based ballots, thus we do not describe them in detail in the further part of this section.

The rules that we are interested in comparing are Rule X, the greedy approval rule, which is a rule currently used for selecting projects, Phragmén's rule, and the sequential version of PAV (sPAV)[15]. We have limited the experiments to polynomial-time algorithms, since the instances are too large to efficiently run e.g. the non-sequential versions of Phragmén's rule or PAV for them.

Since Phragmén's rule does not extend to cardinal utilities, we assume that this rule is used only for approval ballots. Specifically, if we have a preference profile with cardinal utilities constructed from approval ballots, then we give Phragmén's rule as input the initial approval ballots that were used as the base for constructing cardinal utilities.

In our experiments we apply the greedy approval rule in exactly the same way as it is used by the cities, that is this rule is used separately in each district. In contrast, we run Rule X, Phragmén's rule, and sPAV on instances constructed by merging district-wide elections. This way, we analyse whether the considered rules would make proportional choices, even if the projects were not preassigned to specific districts.

In our analysis we used the following metrics:

Total utility (util). The total utility of the voters from the selected set of projects W, that is $\sum_{i \in N} \sum_{c \in W} u_i(c)$.

Distribution of projects (proj-dis). For each election instance we look at the projects selected from each district. We compute their cost and divide it by the fraction of the budget that is proportional to the population of the district (excluding city-wise projects that have been selected). From those ratios we take a variance.

Distribution of utilities (util-dis). For each election instance and each voter i we compute her normalised utility from the set of selected projects W, which we define as $\sum_{i \in N} \sum_{c \in W} u_i(c)$ divided by $n \cdot \sum_{c \in W} u_i(c)$. We compute the variance of these values.

We normalise each of these three metrics: when computing the value of the metric for a committee W we divide it by the value of the metric for the committee returned by the greedy approval rule.

E.1 Comparing Exhaustive Variants of Rule X

We start by comparing three different ways of making Rule X exhaustive. The first two methods, Exh1 and Exh2 were described in Appendix B.3. The third method, Exh3, uses the utilitarian strategy: we first select projects using Rule X, and then add projects greedily,

[15] In fact, instead of PAV we use a rule that greedily maximizes smoothed Nash welfare, i.e., a rule that in every round picks a candidate c maximizing $\sum_{i \in N} \log(1 + u_i(W \cup \{c\}))$, where W is a committee elected so far. In such form, the rule has a straightforward generalization to cardinal utilites.

Election	**util**			**util-dis**			**proj-dis (var)**		
	Exh1	**Exh2**	**Exh3**	**Exh1**	**Exh2**	**Exh3**	**Exh1**	**Exh2**	**Exh3**
Approval utilities									
Cracow-18	1.28	1.39	1.68	0.67	0.64	0.58	1.61	1.39	11.11
Cracow-19	1.33	1.61	1.9	0.8	0.66	0.62	2.71	1.86	17.57
Cracow-20	1.11	1.5	1.68	0.99	0.69	0.66	2.9	2.47	12.35
Czestochowa-20	0.8	1.13	1.16	1.23	0.9	0.92	5.45	5.34	7.09
Warsaw-17	1.26	1.36	1.54	1.11	0.91	0.99	41.77	7.37	26.09
Warsaw-18	1.33	1.43	1.69	1.01	0.9	0.98	34.05	8.73	54.89
Warsaw-19	1.33	1.37	1.63	1.24	1.07	1.25	35.98	6.98	50.48
Cardinal utilities									
Cracow-18	0.25	0.39	1.69	0.42	0.39	0.66	1.64	1.5	58.13
Cracow-19	0.29	0.46	1.8	0.46	0.33	0.62	2.8	1.9	4.58
Cracow-20	0.21	0.61	2.0	0.29	0.28	0.38	2.78	2.98	4.16
Czestochowa-20	0.08	1.33	1.5	0.74	0.76	0.67	4.68	6.46	66.12
Warsaw-17	0.63	1.07	2.19	0.56	0.6	1.69	63.01	13.21	361.12
Warsaw-18	0.8	1.16	2.24	0.72	0.7	1.67	41.28	18.18	432.02
Warsaw-19	0.68	1.07	1.98	0.73	0.66	1.09	55.11	24.81	316.19

Table 1: The results of the experiments comparing exhaustive variants of Rule X.

in each round selecting a project that maximises the ratio of the total utility to the cost and that fits in the remaining budget.

The results of our experiments for approval utilities are presented in Table 1. We make the following conclusions:

1. Different exhaustive variants of Rule X vary significantly. Rule X gives a lot of flexibility to a mechanism designer, selecting often smaller outcomes, yet still satisfying strong fairness requirements, such as EJR. Depending on the specific objectives, a mechanism designer can decide to complete this outcome using different strategies. If the total utility is a primary criterion, then using variant Exh3 gives considerably better results than the other two variants.

2. Both Exh1 and Exh2 produce outcomes that balance voters' satisfaction much better than the current solution (in this sense, both of them are more proportional). However, the outcomes produced by Exh2 are qualitatively better and divide the budget between districts in a substantially fairer way than the ones produced by Exh1. Therefore we suggest Exh2 as the preferred method, when proportionality is the primary criterion.

E.2 Comparing Rule X, Phragmén's rule, and PAV

In our second set of experiments we have compared Rule X with Phragmén's rule, and PAV. The results of those experiments are presented in Table 2. We make the following observations:

Election	util			util-dis (var)			proj-dis (var)		
	RX	**Phrag**	**PAV**	**RX**	**Phrag**	**PAV**	**RX**	**Phrag**	**PAV**
Approval utilities									
Cracow-18	1.39	1.39	0.38	0.64	0.64	1.95	1.39	1.39	19.14
Cracow-19	1.61	1.6	0.64	0.66	0.65	1.36	1.86	1.65	1.15
Cracow-20	1.5	1.5	0.77	0.69	0.7	0.99	2.47	2.49	1.29
Czestochowa-20	1.13	1.1	0.83	0.9	0.96	1.05	5.5	4.41	20.97
Warsaw-17	1.36	1.37	0.4	0.91	0.93	1.19	7.37	6.24	86.22
Warsaw-18	1.43	1.44	0.47	0.9	0.91	1.4	8.73	10.14	102.03
Warsaw-19	1.4	1.41	0.37	1.07	1.1	1.32	6.98	8.59	83.46
Warsaw-20	1.41	1.42	0.91	0.89	0.91	0.96	1.14	1.16	0.24
Warsaw-21	1.62	1.64	0.93	0.73	0.8	0.84	1.06	1.09	0.19
Cardinal utilities									
Cracow-18	0.39	0.28	1.46	0.39	0.37	0.72	1.49	1.28	33.98
Cracow-19	0.46	0.4	1.55	0.33	0.36	0.81	1.94	1.78	1.11
Cracow-20	0.61	0.59	1.83	0.28	0.29	0.44	3.08	2.33	1.16
Czestochowa-20	1.33	0.47	1.63	0.75	1.15	0.66	6.33	4.28	57.07
Warsaw-17	1.07	0.92	1.43	0.58	0.56	1.45	13.67	12.33	154.05
Warsaw-18	1.17	0.96	1.4	0.7	0.69	1.36	18.86	14.35	118.95
Warsaw-19	1.06	0.84	1.36	0.66	0.71	1.2	24.78	17.08	179.77
Warsaw-20	1.01	0.79	1.28	0.83	0.91	0.89	1.16	1.13	15.09
Warsaw-21	1.19	0.68	1.35	0.58	0.45	0.73	1.13	1.04	10.13

Table 2: The results of the experiments comparing Rule X, Phragmén's rule, and PAV.

1. For approval utilities, Rule X and Phragmén's rule give very similar results. According to the utilitarian criterion, and to the distribution of utilities, these results are better than the results returned by the currently used method. Also, those rules divide the budget proportionally among the districts (the variance in the PROJ-DIS criterion is relatively low). PAV is considerably worse, both in terms of the total utility of the selected projects, and in terms of proportionality.

2. For cardinal utilities we observe a difference between Rule X and Phragmén's rule. This difference is expected since Phragmén's rule does not take into account the more fine-grained information on utilities, but operates only on approval ballots. On the other hand, our results suggest that there is indeed a considerable advantage of using rules that use cardinal utilities. We further observe that for cardinal utilities PAV returns outcomes with a higher total utility, but this happens at the huge cost of proportionality. Rule X always selects proportional outcomes, and outperforms Phragmén's rule with respect to the utilitarian criterion.