Core Stability in Participatory Budgeting: Approximations and Open Questions

KAMESH MUNAGALA, Computer Science Department, Duke University, USA
YIHENG SHEN, Computer Science Department, Duke University, USA

Participatory budgeting is an emerging and significant field where social choice theory can be applied. This process aggregates individual preferences to decide on the allocation of public funds. Stability properties are crucial for assessing the outcomes of participatory budgeting. In this survey, we focus on the stability concept of *core* within the context of participatory budgeting. We briefly summarize the key findings from our recent research on three aspects of this topic: algorithms for finding a constant approximation to the core in participatory budgeting [Munagala et al., 2022b]; auditing mechanisms for core stability [Munagala et al., 2022a]; and core stability under allocation constraints [Mavrov et al., 2023]. Finally, we introduce new directions and open problems that have arisen from these studies.

Acknowledgment: This work is supported by NSF grant CCF-2113798.

Contents

1 INTRODUCTION

Consider a scenario where a city is deciding how to spend all the taxes collected from its residents. There are many different candidate public projects to fund. The project can vary in their costs, and different residents may have different ideas on how to spend the taxes. For instance, families with children may prefer a public school or a children's hospital, while others may prefer a park; residents with different locations are more likely to prefer projects near their living areas. The policymaker needs to aggregate these preferences of the residents. The problem of solving such aggregation and deciding which projects to fund is called *participatory budgeting*. According to the New York City Council, "there are more than 3,000 participatory budgeting processes around the world, most at the municipal level" [PB2, 2023]. This indicates that participatory budgeting plays a critical role in city construction.

In a participatory budgeting problem, we denote the budget limit (i.e. the taxes collected) by b. There is a list of m projects (forming a set C), where each $j \in C$ is associated with a cost s_j. The city needs to select a subset W of these projects whose total cost is at most the budget, that is, $\sum_{j \in W} s_j \leq b$. There are n residents or *voters* (forming a set V) in the city, and each voter i has preferences on how the city should spend its budget, denoted by a utility function u_i on project subsets. The goal is, after all the residents submit their preferences on the projects, to find such a subset W that fairly aggregates the preferences.

1.1 Committee Selection and Multiwinner Elections

In the terminology of social choice, the participatory budgeting problem can be viewed as *the committee selection problem* with weighted candidates. Each project corresponds to a candidate in an election, and each resident corresponds to a voter attending the election. There are m candidates and n voters. Each candidate j has size s_j. The subset of candidates we finally select is called the *committee*, and we want to select a committee W within the committee size constraint b.

When the sizes of all the candidates are identical, this problem is reduced to the multiwinner election problem [Aziz et al., 2019, Brams et al., 2007, Brandt et al., 2016, Chamberlin and Courant, 1983, Endriss, 2017, Monroe, 1995, Thiele, 1895]. This problem is central to the social choice and has attracted attention for over a century. In this problem, there is a set V of n voters and a set C of m candidates, out of which a *committee* of k candidates needs to be chosen. Voters express preferences over subsets of candidates, by *utility functions* $\{u_i\}_{i \in [n]}$ where each u_i is a set function of the selected candidates. There are many specifics of multiwinner elections by adding constraints on the utility functions. For example, the *approval election* is the case where each voter i has an approval set $A_i \subset C$ and $u_i(W) = |A_i \cap W|$.

For simplicity, even in participatory budgeting, we use the committee selection terminologies $\langle$*candidates*, *voters*, *committee*, *sizes*$\rangle$ to represent $\langle$projects, residents, funded projects, costs$\rangle$ respectively in this paper.

1.2 Fairness via Core and Its Multiplicative Approximation

Recall that there are n voters forming a set V, and m candidates forming a set C, where candidate j has size s_j. We need to choose a subset W of candidates with total size at most b (that is, $\sum_{j \in O} s_j \leq b$). Denote the utility of voter i for any committee $T \subseteq C$ by a *utility function* $u_i(T)$. We will basically assume this function is non-negative and monotone, with $u_i(\emptyset) = 0$.

A naive approach to solving the participatory budgeting problem is to maximize the social welfare, that is, select a committee W such that $\sum_{i=1}^{n} u_i(W)$ is maximized. Such a solution is typically considered unreasonable, in the sense that it often under-represents the utility of some minority groups. Consider the following approval election:

Example 1.1. There are 4 unit-size candidates $\{a, b, c, d\}$ and 101 voters. $b = 3$, i.e. we can only choose 3 out of the 4 candidates. The first group of 51 voters have approval set $\{a, b\}$ and the second group of the rest 50 candidates have approval set $\{c, d\}$.

The 2-candidate committee that maximizes $\sum_{i=1}^{n} u_i(W)$ is $\{a, b\}$. However, one can imagine that the second group is unhappy since they occupy nearly half the population, but they do not get even 1 candidate (one-third of the size limit) selected in their approval set. These instances raise the thoughts of including *proportional representation* as a fairness constraint in the selection of committees. In the recent decade, different extensions of proportional representation have been developed and widely studied [Aziz et al., 2017, 2018, Brill et al., 2017, Fernández et al., 2017, Monroe, 1995, Sánchez-Fernández et al., 2017, Skowron et al., 2015].

While there are copious notions of fairness for committee selection, the *core* is a classic and influential one among them. This idea has existed for more than a century [Droop, 1881, Lindahl, 1958, Thiele, 1895], and serves as one of the strongest notions of proportional representation. Towards defining this concept, imagine we split the size b among all the voters so that each voter has an *endowment* of $\frac{b}{n}$ that they can use to "buy" candidates. A candidate of size s_j requires an endowment of s_j to "buy". A committee $W \subseteq C$ with total size at most b is said to be in the core, if no subset S of voters can deviate and purchase another committee $T \subseteq C$ by pooling their endowments, so that each voter in S prefers the new committee T to the original one W. Note that the total endowment of S is $|S| \cdot \frac{b}{n}$, so that this set of voters can buy a committee T of size at most $|S| \cdot \frac{b}{n}$. Formally,

Definition 1.2 (Core). A committee W is in the *core* if there is no $S \subseteq V$ and committee $T \subseteq C$ with $\sum_{j \in T} s_j \leq \frac{|S|}{n} \cdot b$, such that $u_i(T) > u_i(W)$ for every $i \in S$.

The core has a "fair taxation" interpretation [Foley, 1970, Lindahl, 1958]. The quantity $\frac{b}{n}$ can be thought of as the tax contribution of a voter, and a committee in the core has the property that no sub-group of voters could have spent their share of tax money in a way that *all* of them were better off. As such it subsumes notions of fairness such as Pareto-optimality, proportionality, and various forms of justified representation [Aziz et al., 2017, 2018, Fernández et al., 2017] that have been extensively studied in multiwinner election and fairness literature.

Despite the satisfying properties of the core, its strength is also its limitation: Even in the simple setting of unit sizes, integer budget, and additive utilities (the so-called *approval-set* setting with general utilities), the core can be empty. (See for example, [Fain et al., 2018].)

A natural approach to circumvent this problem is to show the existence of a committee that *multiplicatively* approximates the core. We define the α-core as follows.

Definition 1.3 (α-Core). A committee W is in the *α-core* if there is no $S \subseteq V$ and $T \subseteq C$ with $\sum_{j \in T} s_j \leq \frac{|S|}{n} \cdot b$, such that $u_i(T) > \alpha \cdot u_i(W \cup \{q\})$ for every $i \in S$ and $q \in C$. We call the candidate q an *additament.*

Note that we introduce the additament in the definition, since no multiplicative approximation is possible without it even in the setting with unit candidate sizes and additive utilities. This follows from examples in previous work [Cheng et al., 2020, Fain et al., 2018]. The idea of a bicriteria (multiplicative and additive) approximation to utilities was first presented in [Fain et al., 2018]. The work of [Peters et al., 2021] presents an almost identical definition as Definition 1.3, except that the additament q must come from the set T. This makes their definition more restrictive and the α-core smaller. They show that when utilities are additive, an $O\left(\log \frac{u_{\max}}{u_{\min}}\right)$-core solution not only exists, but can also be computed in polynomial time, where $u_{\max}$ and $u_{\min}$ are the largest and smallest non-zero utilities any voter has for any feasible committee.

1.3 Utility Functions

Each voter i is associated with a non-negative function $u_i(\cdot)$, where $u_i(T)$ captures their utility for committee $T \subseteq C$. We assume these functions satisfy two properties:

- *Monotonicity.* $u_i(T) \leq u_i(T \cup \{j\})$ for all $T \subseteq C$ and $j \in A$, with $u_i(\emptyset) = 0$.
- 1-*Lipschitz.* $u_i(T) - u_i(T \setminus \{j\}) \leq 1$ for all $T \subseteq C$ and $j \in A$.

Note that the core is *scale-invariant*, so that the definition is robust to scaling utility functions differently for different voters. Therefore, the 1-Lipschitz condition on the utilities is w.l.o.g. If there are no other constraints on the utilities, we call them *general*. In this paper, we will consider several natural utility functions in increasing order of generality:

- Approval. Each voter i has an *approval set* $A_i \subseteq C$. Their utility for T is $u_i(T) = |T \cap A_i|$.
- Additive. Each voter i has utility u_{ij} for $j \in C$. For committee T, $u_i(T) = \sum_{j \in T} u_{ij}$.
- Submodular. For any $T_1 \subseteq T_2$ and $j \in T_1$:

$$u_i(T_1) - u_i(T_1 \setminus \{j\}) \geq u_i(T_2) - u_i(T_2 \setminus \{j\}).$$

- XOS [Feige, 2006, Lehmann et al., 2001]: For additive functions $\{u_{ijq}, j \in C, q \in [\ell]\}$,

$$u_i(T) = \max_{q=1}^{\ell} \sum_{j \in T} u_{ijq}.$$

- β-self bounding [Boucheron et al., 2009]. Given constant $\beta \geq 1$, for each $T \subseteq C$:

$$\sum_{j \in T} (u_i(T) - u_i(T \setminus \{j\})) \leq \beta \cdot u_i(T).$$

We note that approval utilities are a special case of additive, which are a special case of submodular, which are a special case of XOS, which are a special case of 1-self bounding [Boucheron et al., 2009]. Note that though XOS functions are sub-additive, in general, β-self bounding functions *need not be sub-additive*, where sub-additivity means that $u_i(A \cup B) \leq u_i(A) + u_i(B)$ for all $A, B \subseteq C$.

To motivate these classes, approval utilities capture the classical setting of "approval ballots" in elections, and have a rich history in social choice. See the recent book [Lackner and Skowron, 2023] for a comprehensive survey of this topic. Submodular functions capture diminishing returns from choosing additional candidates, and have been widely studied as a discrete analog of concavity.

XOS functions can be motivated in settings where individuals vote on *behalf of a family*. Consider Participatory Budgeting, where the projects either pertain to children or adults, and are additive within each group. An individual voting on behalf of themselves and their children may feel their taxes have been well spent if the *maximum* utility received by anyone in their family is large.

Similarly, in graph theory, the maximum size of a subgraph for any hereditary property is XOS (see [Dubhashi and Panconesi, 2009]). Such functions can capture *diversity or harmony* in the committee. Consider approval utilities with a twist: There is a graph G on candidates, where an edge captures "too similar", say in terms of opinion. Given committee W and voter i's approval set A_i, their utility is the maximum independent set of the sub-graph induced on $W \cap A_i$. This captures opinion diversity in the subset of approved candidates that are on the committee, and is XOS since independent set is hereditary. On the other hand, if the graph edges model a social network and are interpreted as "gets along with", the voter's utility may be the maximum size of a clique in $W \cap A_i$, which corresponds to the maximum sub-committee among approved candidates that all get along. This captures "harmony" in the committee from the voter's perspective, and is XOS as well.

If instead of defining the utility from diversity (resp. harmony) as the size of the maximum independent set (resp. max clique), this is defined as $u_i(W) = \log N(A_i \cap W)$, where $N(A_i \cap W)$ is the number of independent sets (resp. cliques) in the subgraph on $W \cap A_i$, such utilities are called "combinatorial entropies" and remain 1-self bounding [Boucheron et al., 2009].

2 ALGORITHMS FOR CONSTANT APPROXIMATION TO THE CORE

In this section, we summarize the first line of work – how we achieve a constant approximation to the core under different utility functions. In Section 2.1, we introduce and analyze some basic committee selection rules. In Section 2.2, we list our results on the constant approximate core solution. Moreover, we briefly discuss how the basic committee selection rules are related to our algorithms, and how we extend these rules for more general settings.

2.1 Proportional Approval Voting and Its Generalizations

Proportional Approval Voting (PAV). The PAV rule is a classical committee selection rule for multiwinner elections with APPROVAL utilities, dating back a century to Thiele [Thiele, 1895]. This rule plays an important role in our analysis and is closely related to the core stability. For integer $x \geq 1$, let $H(x) = \sum_{y=1}^{x} \frac{1}{y}$ denote the harmonic sum till x. We define $H(0) = 0$. The PAV score of a committee W is defined as:

$$\mathbf{pav}(W) = \sum_{i=1}^{n} H(u_i(W)). \tag{1}$$

Consider the following algorithm that we will term LOCAL:

> LOCAL. Given the current committee W of size k, if there is a $j_1 \in W$ and $j_2 \notin W$ such that $\mathbf{pav}(W \cup \{j_2\} \setminus \{j_1\}) > \mathbf{pav}(W)$, then replace W by $W \cup \{j_2\} \setminus \{j_1\}$.

When this process terminates, we have a local optimum for the **pav** score. The work of [Aziz et al., 2017, Fernández et al., 2017] shows that any such local optimum satisfies a special case of the core termed *extended justified representation* (EJR), where the blocking coalitions satisfy certain cohesiveness conditions. More recently and more relevant to us, it was shown by [Peters and Skowron, 2020] that any such local optimum also lies in the 2-approximate core:

THEOREM 2.1 (2-CORE FOR APPROVAL UTILITIES [PETERS AND SKOWRON, 2020]). *For approval utilities,* LOCAL *finds a committee in the* 2*-approximate core.*

Further, they show this result is tight – any rule that maximizes the sum of symmetric concave functions over voters' utilities cannot do better than a 2-approximation. (As an aside, it is an open question whether a 1-approximate core exists for this setting via a rule not based on scoring functions.)

Generalizations of PAV. In this paper, we will consider modifications of the PAV rule to allow for real-valued utility functions. We first define *Smooth Nash Welfare*, which has been previously studied in [Fain et al., 2018, Fluschnik et al., 2019]. The score of committee W is defined as:

$$\mathbf{snw}(W) = \sum_{i=1}^{n} \ln(1 + u_i(W)). \tag{2}$$

The second generalization is new, and we term it *Generalized PAV*. For $x \geq 0$, we define

$$\Phi(x) = H(\lfloor x \rfloor) + \frac{x - \lfloor x \rfloor}{\lceil x \rceil}.$$

Then the score of committee W is defined as:

$$\mathbf{gpav}(W) = \sum_{i=1}^{n} \Phi(u_i(W)). \tag{3}$$

These rules are very similar to each other. The **gpav** rule reduces to PAV for approval utilities, and satisfies properties like EJR there. On the other hand, the **snw** rule is analytically simpler and leads to somewhat better approximation bounds in our analysis.

The argument in [Peters and Skowron, 2020] can be extended to show that a local optimum for **snw** lies in the 2-approximate core with submodular utilities. However, submodular utilities represents the limit to which Local lies in the approximate core. Once we consider very simple XOS utilities, the following example that local optima to **snw** or **gpav** need not lie in any γ-approximate core for constant γ.

Example 2.2. There are $m = 2k$ candidates and $n = k$ voters, where k is the committee size. There are two sets of k candidates each: $A = \{a_1, \ldots, a_k\}$, and $B = \{b_1, \ldots, b_k\}$. The utility function of voter i is as follows: For set T, $u_i(T) = \max(|T \cap B|, |T \cap \{a_i\}|)$.

Since the utility function u_i is the maximum of two additive functions, it is XOS. Consider the committee $W = A$. If any a_i is replaced by any b_j, the utilities of all voters are unchanged at value 1. Therefore, $W = A$ is a local optimum to **snw** (resp. **gpav**). However, all voters can together choose blocking committee B, which gives each of them a factor k larger utility. Therefore, the local optimum A does not lie in the γ-core for any constant γ.

Improved Analysis of **pav** *for Large Coalitions.* In the multiwinner election problem with additive utilities, if we use the Local rule with the **pav** score, Peters and Skowron [2020] show an approximation factor of 2, which is tight. However, this tightness holds only for *small coalitions* of voters. This begs the question: *Is there an improved analysis of the* Local *rule for any coalition size?* We answer this in the affirmative: We show that as the coalition size increases, the approximation factor of the local optimum to **gpav** approaches 1. In particular, this shows Local is weakly Pareto-optimal. In addition, our analysis holds for general additive utilities (and not just approval), which shows the desirability of **gpav** as a scoring rule.

Theorem 2.3. *For multiwinner elections with additive utilities (and no allocation constraints), suppose only coalitions of size at least αn are allowed to deviate, where $\alpha \in [0, 1]$. Then any local optimum to* **gpav** *lies in the $2 - \alpha$ approximate core. Further, this bound is tight for such local optima.*

2.2 Finding a Constant Approximate Core Solution in Participatory Budgeting

We show the existence of an $O(1)$-core in Participatory Budgeting when the utility functions of the voters are monotone and submodular in [Munagala et al., 2022b]. This is an improvement over the work of Peters et al. [2021] mentioned above, which presents a logarithmic approximation for the restricted case of additive utilities.

Theorem 2.4 (PB, Submodular Utilities, Polytime). *For monotone submodular utilities, a 67.37-core is always non-empty. One such solution can be computed in polynomial time.*

We also improve the constant for the well-studied special case of additive utilities:

Theorem 2.5 (PB, Additive Utilities, existence). *For additive utilities, a 9.27-core is always non-empty.*

Unlike Theorem 2.4, we do not know how to implement the algorithm in Theorem 2.5 in polynomial time. We remark that the previous two results can be combined to show that a 15.2-core solution for additive utilities can be computed in polynomial time.

Theorem 2.6 (PB, Additive Utilities, Polytime). *For additive utilities, a 15.2-core is always can be computed in polynomial time.*

Our overall algorithm procedure is to apply the multilinear extension to extend the utility function to fractional committees. Then we use a continuous time local search procedure similar to Local for the Nash Welfare objective to find a fractional solution, such that it almost lies in the 2-approximate fractional core. We subsequently round it iteratively to find a solution in the approximate integer core.

Lower Bounds. The next natural question is whether our results can be extended to arbitrary monotone utilities. We show that this is not possible. We present an example to show that an $O(1)$-core (or even any $f(n,m)$-core where n is the number of voters and m is the number of candidates) may not exist for general monotone utilities.

Theorem 2.7 (General Utilities, Lower Bound). *For general monotone utilities, for any function $\varphi : \mathbb{Z}^+ \times \mathbb{Z}^+ \to \mathbb{R}^+$, a $\varphi(n,m)$-core can be empty.*

We also show that the case of submodular utilities does need a multiplicative approximation to the core, and the 1-core can be empty. This justifies the form of Theorem 2.4 that involves both a multiplicative approximation and an additament.

Theorem 2.8 (Submodular Utilities, Lower Bound). *For monotone submodular utilities, a 1.015-core can be empty.*

In Mavrov et al. [2023], we relax restriction on utility function from submodular to more general functions. We present an existence proof of a $e^{O(\beta)}$-core for β-self bounding utilities, after slightly weakening the definition of the additament.

Theorem 2.9 (β-Self bounding Utilities). *For the Participatory Budgeting problem with β-self bounding utilities (where $\beta \geq 1$ is an integer) and no allocation constraints, a c-approximate core is always non-empty, where $c = e^{O(\beta)}$.*

Finally, we complement this by showing that the exponential dependence of the approximation factor on β in Theorem 2.9 is unavoidable:

Theorem 2.10 (β-Self bounding Utilities, Lower Bound). *For multiwinner elections with β-self bounding functions ($\beta \geq 5$) and no allocation constraints, the c-approximate core can be empty for $c = \frac{1}{2}\left(\frac{4}{3}\right)^{\beta/2} - o(1)$.*

3 APPROXIMATION ON ENDOWMENT AND AUDITING

In this section, we first present an alternate definition of approximate core stability and introduce the problem of committee auditing. Next, in Section 3.2, we briefly list our results for auditing the core and other stability notions inspired by core auditing.

Approximate Endowment Core. Since the statement in the core definition is composite, ways of relaxing the definition can vary. Other than the previous approximate notion of α-core, people have defined an alternate endowment version of approximation. Intuitively, instead of relaxing the definition by forcing the deviating subset to get more utility, the endowment approximation restricts their "power" to form a new committee. The approximation ratio reflects how much the budget shrinks if a voter subset is going to deviate from the current committee. The definition is as follows:

Definition 3.1 (θ-Endowment Core). For $\theta \leq 1$, a committee W of size at most k lies in the *θ-endowment core* if for all $S \subseteq [n]$, there is no *deviating committee* T with size at most $\theta \cdot |S| \cdot \frac{k}{n}$, such that for all $i \in S$, we have $U_i(T) > U_i(W)$.

It is known [Jiang et al., 2020] that a $\frac{1}{32}$-endowment core solution always exists for very general utility functions of the voters.

3.1 Auditing for Endowment Core Stability

Though the existence of approximate endowment core solutions is a strong positive result, the algorithms for finding such solutions are often complex. Indeed, even in settings where the core is known to be always non-empty, for instance when candidates can be chosen fractionally [Foley, 1970], the non-emptiness is an existence result that needs an expensive fixed point computation. On the other hand, in practice, what are implemented are typically the simplest and most explainable social choice methods such as Single Transferable Vote (STV). Therefore, from the perspective of a societal decision maker, such as a civic body running a participatory budgeting election, it becomes important to answer the following *auditing* question for any given election:

> Given a committee W of size at most k found by some implemented preference aggregation method, how close is it to being core stable, i.e., what is the smallest value of θ_c such that W *does not* lie in the θ_c-endowment core for that instance?

Note that if a committee W lies in the core, then $\theta_c > 1$, else $\theta_c \leq 1$. Such an auditing question is useful even if the decision maker themselves is not sensitive to fairness because it allows for the review of implemented decision rules via a third party or government agency. Further, the set of deviating voters that correspond to the θ_c-approximation yields a demographic that is unhappy with the current outcome, and this can be analyzed further by policy makers. We term the above question as the *core auditing* problem.

3.2 Hardness and Approximation Algorithm for Core Auditing

In [Munagala et al., 2022a], we show that for Approval Elections, the value of θ_c in the core auditing problem is NP-Hard to approximate to a constant factor:

Theorem 3.2 (Hardness of Auditing). *For any constant $\gamma > 0$, approximating θ_c to within a factor of $1 + \frac{1}{e} - \gamma$ is* NP-Hard.

Furthermore, we show that such APX-Hardness persists even when voters are allowed to choose a fractional deviating committee, with a proper definition of the approximate core stability for fractional committees.

Definition 3.3. For $\theta \leq 1$ and constant $\eta \in (0, 1]$, a committee $\vec{x} \in [0, 1]^m$ with $\sum_j s_j x_j \leq k$ lies in the (θ, η)*-approximate fractional core* if for all $S \subseteq [n]$, there is no *deviating committee* $\vec{y} \in [0, 1]^m$ with $\sum_j s_j y_j \leq \theta \cdot |S| \cdot \frac{k}{n}$, such that for all $i \in S$, we have $U_i(\vec{y}) \geq U_i(\vec{x}) + \eta$.

Theorem 3.4 (Hardness of Auditing for fractional committee). *For any $0 < \eta \leq 1$ and any $\gamma > 0$, distinguishing instances that do not lie in the (θ_c, η)-approximate fractional core from those that lie in the $(\theta_c(1.1839 - \gamma), \eta)$-approximate fractional core is* NP-Hard.

On the positive side, we design an algorithm for computing an $O(\min(\log m, \log n))$ approximation for the value θ_c of a given committee:

Theorem 3.5. *Given a committee W of size at most k, its θ_c value can be computed within $O(\min(\log m, \log n))$ factor in polynomial time, where m, n are the total number of candidates and voters respectively.*

The idea of the algorithm is via linear program rounding. Our program (and indeed, our auditing question itself) is an interesting generalization of the densest subgraph problem [Charikar, 2000], where the goal is to choose a subgraph with maximum average degree. Given a graph, treat voters as edges and candidates as vertices that are approved by the incident edges; further assume any voter needs utility 2 (that is, both end-points) in a feasible deviation. Then, the value of θ_c is precisely the density of the densest subgraph (to scaling). We combine ideas from the rounding for densest subgraph (where the rounding produces the integer optimum without approximation) with that from maximum coverage to design our rounding scheme. We further show that our linear program has an integrality gap of $\Omega(\min(\log m, \log n))$, showing that we cannot do any better against an LP lower bound. We also extend this to general candidate sizes and arbitrary additive utilities via knapsack cover inequalities, leading to an $O(\min(\log m, \log n))$ approximation factor in participatory budgeting core auditing.

Deciding Core Stability. A problem closely related to the auditing problem is to decide if a committee W does not lie in the core – this is equivalent to deciding whether its $\theta_c \leq 1$. The problem is known to be NP-hard in Brill et al. [2020]. We show that this decision problem is NP-Hard even in a "constant degree" setting:

Theorem 3.6 (Hardness of Deciding the Core, Constant Approval Set Size). *Deciding whether a committee W does not lie in the core (that is, deciding whether its $\theta_c \leq 1$) is* NP-Hard *when each voter approves at most 6 candidates (that is, $|A_i| \leq 6$ for all voters $i \in [n]$), and each candidate lies in at most 2 of the sets A_i.*

4 ELECTIONS WITH ALLOCATION CONSTRAINTS

Observing the common fact that real-world elections have restrictions on the selected committee, we further focus on the practically relevant aspect of having exogenous constraints on a feasible committee. We assume there is a set $\mathcal{P}$ of feasible committees (each of size at most k), and the chosen committee W must belong to this set. For simplicity, we only focus on unit-size candidates, i.e., the committee selection problem.

Several types of constraints could arise in practice, and we now give some examples.

- *Matroid Constraint.* Multiwinner elections with a single matroid constraint were previously considered in [Fain et al., 2018]. Here, $\mathcal{P}$ consists of all independent sets of size at most k in the matroid $\mathcal{M}$. The simplest example of matroids is a partition matroid constraint. The set C of candidates are partitioned into disjoint groups $G_1, G_2, \ldots, G_\ell$, and any feasible committee of size k can choose at most k_i candidates from group G_i, where the k_i are exogenously specified. As an example, the groups could represent geographic regions the candidates hail from, or the type of project in Participatory Budgeting.
- *Packing Constraints.* Here, there are multiple downward-closed constraints, meaning that any sub-committee of a feasible committee is also feasible. For instance, imagine candidates belong to multiple overlapping groups (different races, genders, income levels), and there is a constraint on the number of candidates that can be chosen from any group.
- *Independent Set.* This is a special case of packing constraints. We have a graph over the candidates, with the constraint that a feasible committee is an independent set in this graph. This captures pairs of candidates who have conflicts or pairs of projects that cannot simultaneously be funded. These projects cannot be simultaneously put on the committee.
- *Rooney Rule.* Going beyond packing constraints, we can have minimum (or covering) requirements. For instance, if we seek diversity in the selected candidates, we could impose minimum numbers on candidates chosen from certain groups. As an example, a committee needs to include at least x female candidates, or a Participatory Budgeting outcome needs to include at least one public safety project and at least two child-friendly projects.

We consider the most general model where the set $\mathcal{P}$ of feasible committees of size at most k can be an *arbitrary subset* of 2^C. Though it is tempting to use Definition 1.3 while restricting the blocking committee T to also lie within $\mathcal{P}$, the γ-core may be empty for any constant γ even for a single packing (or partition matroid) constraint.

4.1 Restrained Core and Its Multiplicative Approximation

We use the perspective of social planner protecting the rights of the voters who do not deviate by providing them first their "fair share" of the budget. This leads to our first contribution, defining the *restrained core*. To understand this definition, given allocation $W \in \mathcal{P}$, suppose subset S of voters deviates with its endowment $k' = \lfloor \frac{|S|}{n} k \rfloor$. Then, $S' = V \setminus S$ is also entitled to $k - k'$ candidates. The social planner picks at most $k - k'$ candidates from the current allocation W for S'. This leaves space for S to pick k' candidates from C subject to the feasibility constraint. See Fig. 1 for an illustration of the deviation process. Formally, we define the restrained core as follows:

Definition 4.1 (γ-approximate restrained core). Given a set $\mathcal{P}$ of committees of size at most k, a committee $\hat{W}$ is said to be *q-completable* if there exists W'' with $|W''| \leq q$ such that $W'' \cup \hat{W} \in \mathcal{P}$.

A committee $W \in \mathcal{P}$ lies in the γ-approximate restrained core if there is no constraint-feasible γ-blocking coalition $S \subseteq V$ of voters. Such a blocking coalition with endowment $k' = \lfloor \frac{|S|}{n} k \rfloor$ satisfies the following: For all k'-completable committees $\hat{W} \subseteq W$ with $|\hat{W}| \leq k - k'$, there exists W' with $|W'| \leq k'$ such that (1) $T = W' \cup \hat{W} \in \mathcal{P}$, and (2) for all $i \in S$, it holds that $u_i(T) \geq \gamma \cdot (u_i(W) + 1)$.

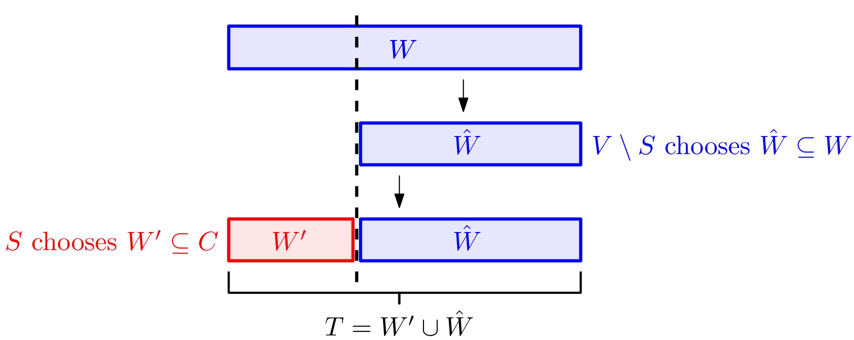

Fig. 1. Illustration of the deviation process in the restrained core. The original committee is shown as W in the first row. Assume that subset S of voters are deviating. The rest of the voters $V \setminus S$ chooses $\hat{W}$ from W with size at most $k - k'$ (the second row). Then S chooses W' with at most k' candidates (shown as red in the third row) from C. The final deviated committee is $T = W' \cup \hat{W}$.

We insist $\hat{W}$ is k'-completable in order to ensure there is always some choice of W' for Condition (1), which is important to make sure the condition is not vacuously false when $|\hat{W}| < k - k'$. Further, note that when $\mathcal{P}$ is the set of all committees of size at most k, that is, when there are no allocation constraints, then Definition 4.1 reduces to Definition 1.3. To see this, simply note that the choice of W' in Definition 4.1 is now not affected by the choice of $\hat{W}$, so that $\hat{W} = \emptyset$ without loss of generality. Therefore, Definition 4.1 generalizes Definition 1.3 to constraints.

4.2 Finding an Approximate Restrained Core

Building on the definition of approximate restrained core, our main technical contribution is the following theorem.

Theorem 4.2 (Approximate Restrained Core). *For multiwinner elections with* arbitrary *allocation constraints $\mathcal{P}$ and β-self bounding utility functions for $\beta \geq 1$, an e^{β}-approximate restrained core is always non-empty. As a consequence, the e^{β}-approximate core is non-empty without allocation constraints.*

As we mention in Section 6, though there has been prior work on core with constraints, these either require scaling down the constraints on deviation often rendering them meaningless, or work in very limited settings. Our Definition 4.1 and the associated Theorem 4.2 are the first results that achieve a constant approximate core for *arbitrary* constraints even for approval utilities.

Since XOS utilities are 1-self bounding [Vondrák, 2010], Theorem 4.2 implies an e-approximate restrained core for XOS utilities (and hence, for approval, additive, and submodular utilities) with any allocation constraints, or an e-approximate core without allocation constraints (Definition 1.3).

One choice of $\hat{W}$ in Definition 4.1 that yields Theorem 4.2 is to maximize the **snw** score for voters not in the deviating coalition. Therefore, the social planner takes care of the complement in the best possible fashion for any deviation, which itself can be viewed as a form of fairness.

Finally, the exponential dependence of the approximation on β is unavoidable; see Theorem 2.10.

Algorithm. The algorithm that yields the above result is surprisingly simple:

> Global: Find $W \in \mathcal{P}$ such that $\mathbf{snw}(W)$ is maximized.

Note that we are not finding a local optimum, but instead computing the *global optimum* of **snw**; indeed, when $\mathcal{P}$ is arbitrary, the Local algorithm may get stuck simply for lack of swaps that preserve membership in $\mathcal{P}$. Further, Example 2.2 shows Local is insufficient for XOS functions even without any additional constraints. Our use of the global optimum necessitates an entirely new analysis compared to prior work, and this analysis forms a key contribution.

We have therefore presented the first fairness analysis of Nash Welfare for multiwinner elections with XOS utilities even without additional constraints. We note that compared to prior work on welfare maximization with XOS utilities [Feige, 2006, Lehmann et al., 2001] that were based on linear programming, our proof for **snw** is entirely combinatorial. This is because we only use the self-bounding property of these functions, while welfare maximization uses the stronger property of fractional subadditivity of XOS functions. To highlight the difference, our results hold for arbitrary self-bounding functions, while welfare maximization results extend to sub-additive functions. These classes are incomparable, and we do not know how to extend our results to sub-additive functions.

Finally, we note that for one voter, core stability reduces to utility maximization, which cannot be approximated in polynomial time within sub-polynomial factors for either XOS functions (value oracle model; [Mirrokni et al., 2008]) or independent set constraints (NP-Hardness; [Feige et al., 1996]). Our results therefore show fairness properties for Nash Welfare even in settings where there are no computationally efficient and fair algorithms possible via any method.

Lower Bound for Restrained Core. One may wonder if Definition 4.1 makes the problem "too easy" so that there is always a 1-approximate (exact) restrained core. We show this is not the case even in the presence of very simple constraints and approval utilities, via the following theorem:

Theorem 4.3 (Restrained Core, Lower Bound). *For $c = 16/15 - o(1)$, a c-approximate restrained core can be empty even for approval utilities and a single packing or partition matroid constraint.*

This lower bound complements the upper bound of e for additive utilities ($\beta = 1$) in Theorem 4.2. Note that in the absence of constraints, it is a long-standing open question whether a 1-approximate (exact) core exists for approval utilities. The above theorem shows that surprisingly, even with a single constraint, the exact (restrained) core for this setting is empty. Indeed, the theorem holds even for a weaker version of Definition 4.1, where $\hat{W}$ could be any committee of size $k - k'$ (and not necessarily a subset of W) such that there exists W' making $\hat{W} \cup W' \in \mathcal{P}$.

4.3 Matroid Constraints

Consider the special case where $\mathcal{P}$ is the set of independent sets of a matroid. We show that the Local rule applied to **snw** lies in the 2-approximate restrained core for a matroid constraint and submodular utilities. In this setting, the Local rule swaps a candidate $j \notin W$ for $\ell \in W$ as long as the committee remains a basis of the matroid and the **snw** score strictly improves. Note that unlike Global, this algorithm is computationally efficient. The proof is built on [Peters and Skowron, 2020], who show a 2-approximate core for the special case of **pav** with approval utilities and no constraints. They also show that the factor of 2 is tight for the **pav** rule without constraints, and the same tightness will hold for our setting.

4.4 Other Fairness notions with allocation constraints

We next consider the notion of *extended justified representation (EJR)* [Aziz et al., 2017], which is a weakening of the core for approval utilities. This is exactly satisfied by Local applied to **pav** rule in the absence of constraints. We define a generalization of EJR to the setting with constraints similar to the way to define restrained core. We call this definition *restrained EJR*:

Definition 4.4 (Restrained EJR for Approval Utilities). We are given a set $\mathcal{P}$ of feasible committees of size at most k. A committee $W \in \mathcal{P}$ satisfied restrained-EJR if there is no constraint-feasible blocking coalition $S \subseteq V$ of voters. Such a blocking coalition with endowment $k' = \lfloor \frac{|S|}{n} k \rfloor$ satisfies the following: For all k'-completable committees $\hat{W} \subseteq W$ with $|\hat{W}| \leq k - k'$, there exists W' with $|W'| \leq k'$ such that

(1) $T = W' \cup \hat{W} \in \mathcal{P}$, and
(2) For all $i \in S$, $|(\cap_{i \in S} A_i) \cap T| \geq \max_{i \in S} u_i(W) + 1$.

We show that **pav** satisfies exact restrained EJR for approval utilities when the constraints form the independent sets of a matroid. In contrast, the exact restrained core for this setting can be empty from Theorem 4.3. More recently, Masařík et al. [2023] proved that the restrained EJR always exists when the constraint $\mathcal{P}$ is *downward-closed*, and it is incompatible with Pareto-optimality. This implies that **pav** can fail restrained EJR when $\mathcal{P}$ is not a matroid.

5 SUMMARY OF APPROXIMATION RESULTS

In Table 1, we present a summary of the results for approximate core under various utility functions, candidate sizes (unit vs. general), and allocation constraints.

6 RELATED WORKS

Proportionality and the Core. One classic objective in committee selection is achieving fairness via proportionality, where different demographic slices of voters feel they have been represented fairly. This general idea dates back more than a century [Droop, 1881], and has recently received significant attention [Aziz et al., 2017, 2018, Brams et al., 2007, Brill and Peters, 2023, Chamberlin and Courant, 1983, Fernández et al., 2017, Monroe, 1995]. In fact, there are several elections, both at a group level and a national level, that attempt to find committees (or parliaments) that provide approximately proportional representation. For instance, the popular Single Transferable Vote (STV) rule is used in parliamentary elections in Ireland and Australia, and in several municipal elections in the USA. This rule attempts to find a proportional solution.

A long line of recent literature has studied the complexity and axiomatization of voting rules that achieve proportionality; see [Aziz et al., 2019, Brandt et al., 2016, Endriss, 2017, Lackner and Skowron, 2023] for recent surveys. Proportionality in committee selection arises in many other applications outside of social choice as well. For example, consider a shared cache for data items in a multi-tenant

Utility	Sizes	Constraints	Approx.	Lower Bd.	Run Time
β-self bounding	General	None	$e^{O(\beta)}$ [2023]	$\frac{1}{2}\left(\frac{4}{3}\right)^{\beta/2}$ [2023]	–
Submodular	General	None	67.37 [2022b]	1.015 [2022b]	Poly.
Additive	General	None	9.27 [2022b]	Unknown.	–
Approval	General	None			–
Additive	General	None	15.2 [2022b]		Poly.
Approval	General	None			Poly.
β-self bounding	Unit	General	e^{β} [2023]	$\frac{1}{2}\left(\frac{4}{3}\right)^{\beta/2}$ [2023]	–
XOS	Unit	General	e [2023]	16/15 [2023]	–
Submodular	Unit	Matroid	2 [2023]		Poly.
Approval	Unit	Matroid			Poly.

Table 1. Summary of results for approximate core. The upper bound of 2 for submodular utilities also holds for approval utilities, while the lower bound of 16/15 for approval utilities also holds for submodular and XOS utilities. An empty box in "Run Time" implies an existence result.

cloud system, where each data item is used by several long-running applications [Friedman et al., 2019, Kunjir et al., 2017]. Each data item can be treated as a candidate, and each application as a voter whose utility for an item corresponds to the speedup obtained by caching that item. In this context, a desirable caching policy provides proportional speedup to all applications. More recently, proportionality fairness has been introduced and studied in sequential decision processes Chandak et al. [2024] and with representatives in metric spaces Kalayci et al. [2024].

The core represents the ultimate form of proportionality: Every demographic of voters feel that they have been fairly represented and do not have the incentive to deviate and choose their own committee of proportionally smaller size which gives all of them higher utility. In the typical setting where these demographic slices are not known upfront, the notion of core attempts to be fair to all subsets of voters. The work of [Munagala et al., 2021] formally argues that in certain multiwinner election settings, the core also approximately optimizes simpler diversity measures of the resulting committee. The core is known to be non-empty in multi-winner elections under certain restrictions [Brill et al., 2024, Pierczyński and Skowron, 2022]. However, for committee selection with standard Approval utilities, the non-emptiness of the core remains open (see Section 7.1).

Fisher Markets. Our fractional solutions are superficially related to the Fisher market equilibrium [Arrow and Debreu, 1954, Brainard and Scarf, 2005, Nash, 1950] when divisible items need to be *allocated* to agents, and agents' utilities are additive. For the Fisher market, the optimum Nash Welfare solution finds market clearing prices. However, in a Fisher market, the prices are common to the agents while the allocations are different, while in a Lindahl equilibrium, the prices are per-voter while the allocation (or committee) is common and provides shared utility to all the voters. This is a key difference – the Fisher market has a polynomial time algorithm via convex programming [Eisenberg and Gale, 1959], while no polynomial time algorithm is known for the Lindahl equilibrium even when candidates have unit sizes and voters' utilities are additive (or linear till the maximum size of the candidate). Similarly, though the Nash Welfare solution finds market clearing prices for the Fisher market via strong duality, in the case of public goods, there is no obvious way to interpret the dual of the Nash Welfare solution as market clearing prices. Moreover, for submodular utilities and multilinear extensions, the Nash Welfare objective is no longer a convex program, so that strong duality does not apply.

Approximate Core. In this paper, we have focused on approximating the utility voters obtain on deviating (see Definition 1.3). As mentioned before, this notion first appeared in [Fain et al., 2018], and the notion of a single additament in approximation is due to [Peters et al., 2021]. The latter work present a logarithmic approximation for the special case of additive utilities. The work of [Peters and Skowron, 2020] shows that the well-known Proportional Approval Voting method [Thiele, 1895] achieves a 2-core for the special case where the utilities are additive and candidates are unit size, with each voter having utility either zero or one for each candidate. This algorithm can be viewed as a discrete version of Nash Welfare, and in essence, we can extend this result to the case of submodular utilities and general costs, showing that it yields a 2-core for the *fractional* case of multilinear extension via a polynomial time local search algorithm. The work of [Chen et al., 2019] presents a constant approximation for the K-clustering problem, where the committee is a set of K centers in a metric space, and the cost of a voter is the distance to the closest center. However, these ideas do not extend to the committee selection problem we consider in this paper.

The work of [Jiang et al., 2020] considers a different notion of approximation: Instead of approximating the utility, they approximate the *endowment* that a voter can use to buy the deviating committee. Building on the work of [Cheng et al., 2020], it shows a different fractional relaxation, to which a 2-approximation always exists. They then iteratively round this fractional solution to an integer solution that is a 32-approximation for *all* monotone utility functions. The problem of approximating utilities is very different; indeed, Theorem 2.7 shows we cannot hope to have a similar constant approximation for all utility functions. Nevertheless, we use the idea of iterative rounding from that work, albeit with an entirely different fractional solution and analysis. In effect, we showcase the power of iterative rounding as a unifying framework for finding approximate core solutions, regardless of the notion of approximation.

Rounding Techniques. The notion of multilinear extension and correlation gap has been widely used in stochastic optimization [Agrawal et al., 2010], mechanism design [Yan, 2011], and rounding [Călinescu et al., 2011, Chekuri et al., 2014, Vondrák, 2008]. Typically, it has been used to develop computationally efficient approaches; on the other hand, we demonstrate an application to showing a purely existential result. Similarly, rounding of market clearing solutions have been used to show approximately fair allocations of indivisible goods among agents [Barman et al., 2018a, Cole and Gkatzelis, 2015]. The structure of these problems (common prices but different allocations) is very different from ours (common allocations and different prices), and we need different techniques. Again, in contrast with the resource allocation literature, we need the rounding just to show an existence result as opposed to a computational one.

Auditing for Fairness. The question of auditing has become salient given the increasing democratization of societal decision making, for instance via processes like participatory budgeting. In the context of social choice, there are natural properties that are easy to achieve algorithmically but hard to audit. For instance, checking if an arbitrary outcome is Pareto-optimal is computationally hard [Aziz et al., 2016], while achieving it via some algorithm is easy. We take a further step in this direction by studying the approximate audit of arguably the strongest possible group fairness notion, the core, as well as related fairness properties.

Going beyond social choice, the notion of auditing for group fairness has gained prevalence in machine learning. Here, the "voters" are data points, and the "committee" is a classifier. We wish to audit if the classifier provides comparable accuracy for various demographic slices. The work of [Kearns et al., 2018] formulates and presents algorithms for this problem.

Nash Social Welfare. The **snw** objective is closely related to Nash Social Welfare [Arrow and Debreu, 1954, Brainard and Scarf, 2005, Nash, 1950]. This has been widely studied in the allocation of *private goods*, where each participant has an additive utility over the bundle of goods they receive. When goods are divisible, Nash Welfare is the solution to the Fisher market equilibrium [Eisenberg and Gale, 1959]. When goods are indivisible, [Caragiannis et al., 2019] show that a *local optimum* to this objective (where pairs of goods can be swapped between individuals) satisfies approximate envy-freeness (EF1). The global optimum of the Nash Welfare objective satisfies Pareto-optimality as well. We note that this setting, there are pseudo-polynomial time algorithms achieving both properties [Barman et al., 2018b]. In contrast, our paper shows fairness properties for Nash Welfare in settings where no computationally efficient approximations to stability are even possible.

Core with Constraints. Prior work has tried addressing the aspect of constraints via either changing the definition of the core or what an approximation means. We now contrast these with the present work. As mentioned before, the work of [Cheng et al., 2020, Jiang et al., 2020] considers a different approximation notion where the endowment of a coalition is scaled down when they deviate. Their results extend to packing constraints of the form $A\vec{x} \le b$, where $\vec{x}$ is a binary vector representing which candidates are present in the committee. However, for a coalition of size αn, they require the deviating committee $\vec{y}$ satisfy $A\vec{y} \le \alpha b$, that is, they change the constraint set to make it more strict. This may make the constraint on deviation impossible to satisfy − for instance, an independent set constraint is of the form $x_j + x_\ell \le 1$, where $j, \ell \in C$. If we replace the RHS by $\alpha < 1$, this forces both of j, ℓ to not be chosen, so that the only feasible committee for any deviation is empty. In contrast, Definition 4.1 does not change the constraint set, and further, works not just for packing constraints, but for other constraints such as the Rooney Rule.

A different notion of core for multiwinner elections, defined in [Fain et al., 2018], is the following: When a coalition S deviates, they are allowed to choose a committee of size k; however, they need to obtain a factor $\gamma \cdot n/|S|$ factor larger utility on deviation for it to be a γ-approximate core. Like our notion, their notion also allows for constraints. Indeed, they consider the same setting as our matroid constraint except with additive utilities and show that the same Local algorithm yields an approximate core solution in their notion as well. However, the approximation factor becomes super-constant for multiple matroid constraints or for general packing constraints, even with approval utilities. Indeed, they show that the core does not exist to any non-trivial approximation for independent set constraints with approval utilities. In contrast, Definition 4.1 extends smoothly to arbitrary constraints, yielding a e^β-approximate core for very general β-self bounding utilities.

7 OPEN QUESTIONS

In this section, we propose some directions and list some open questions related to core stability in participatory budgeting (and committee selection).

7.1 Exact core existence and Lindahl priceability

Simple as it may look, the following problem is well-known in the social choice community and has been open for years:

Open Question 1. In committee selection with Approval utilities, does there always exist a committee W that lies in the exact (i.e. 1-approximate) core?

In our work of auditing the core stability, we have discovered some interesting concepts related to the exact core stability. We term it *Lindahl Priceability*:

Definition 7.1 (Lindahl Priceability). A committee W of size at most k is Lindahl priceable if there exists a price system $\{p_{ij}\}$ from voters to candidates, such that the following hold:

1. $\forall j \in [m],\ \sum_i p_{ij} \le 1$, and
2. $\forall i \in [n],\ T \subseteq C$, if $|T \cap A_i| \ge |W \cap A_i| + 1$, then $\sum_{j \in T} p_{ij} > k/n$.

Assume that each candidate costs 1. Since the committee size is at most k, we have a public budget k. The budget is then split equally among all the voters, so each voter has a budget of k/n. The first condition above means that for each candidate, the prices from all voters sum up to at most 1, so that each candidate is not "over-paid". The second condition means a voter cannot afford any committee that she strictly prefers to W.

Lindahl priceability can be viewed as an integral version of the *gradient optimality conditions* in the fractional Lindahl equilibrium [Foley, 1970]. Analogous to the fractional Lindahl equilibrium, we have shown in [Munagala et al., 2022a] that the following proposition holds:

Proposition 7.2. *If a committee is Lindahl priceable, it lies in the core.*

However, even though Lindahl priceability implies core stability, we do not know if it is *strictly* stronger than the core. We conjecture that these two notions are the same:

Open Question 2. In committee selection with Approval utilities, if a committee W is in the exact core, is it Lindahl priceable?

We also do not know if there always exists a committee that is Lindahl priceable. Since we believe Lindahl priceability is a more interpretable notion than the core stability, the following problem might be easier to think of than Open Problem 1. Note that if the answer to Open Problem 2 is "YES", the following question is equivalent to Open Question 1.

Open Question 3. In committee selection with Approval utilities, does there always exists a *Lindahl priceable* committee W?

7.2 Relaxation to the exact core via loss from the current committee

In [Munagala et al., 2022a], we have proposed a relaxation of core stability in committee selection, in the sense that a group of voters will only deviate if each of them is strictly better off, while not giving up anyone in her approval set in the current committee. We call this *sub-core.*

Definition 7.3 (Sub-core). A committee W lies in the *sub-core* if there is no $S \subseteq V$ and committee T with $|T| \leq \frac{|S|}{n} \cdot k$, s.t. $A_i \cap W \subsetneq A_i \cap T$ for all $i \in S$.

It is easy to show that sub-core always exists, and can be found by many rules, e.g., Phragmén's rule [Brill et al., 2017].

The sub-core can also be viewed as every voter has zero tolerance for the loss from the current committee when they deviate. If voters are deviating from W, they can only deviate to a new committee T when the approved candidates in T form a strict superset of their approved candidates in W. However, when we try to move one step further from the sub-core to core, we get stuck. Consider the following extended version of Definition 7.3, where every voter can tolerate at most a loss of q approved candidates from the current committee W:

Definition 7.4 (q-tolerant core). A committee W lies in the *q-sub-core* if there is no $S \subseteq V$ and committee T with $|T| \leq \frac{|S|}{n} \cdot k$, s.t. $|A_i \cap T| \geq |A_i \cap W| + 1$ and $|(A_i \cap W) \setminus (A_i \cap T)| \leq q$ for all $i \in S$.

Based on the above definition, we observe that the sub-core is equivalent to the 0-tolerant core and the exact core (Definition 1.2) is equivalent to the ∞-tolerant core. Definition 7.4 becomes stronger as q increases, yet even for $q = 1$, we do not know if the 1-tolerant core is always non-empty. Therefore, we propose the following question as a possible "first step" towards the existence of exact core (Open Question 1):

Open Question 4. (Weaker version of Open Question 1) In committee selection with Approval utilities, does there always exist a committee W that lies in the 1-tolerant core?

7.3 Extending Table 1

As readers may realize, our constant approximation results for the core results are generalized only to XOS and self bounding utility functions. The next utility setting where we do not know how to get a constant approximation for unit-size candidates is the sub-additive utilities.

Open Question 5. For unit-size candidates with sub-additive utilities, does there exist a committee that lies in the c-core for some constant c?

On the other side, since all our restrained core results are for unit-size candidates, a natural question to ask is whether we can extend these approximation results to arbitrary candidate sizes:

Open Question 6. Can we extend Definition 4.1 to participatory budgeting and also get a constant approximation to the restrained core in participatory budgeting with general candidate sizes?

7.4 Different notions of Core Approximations

In this paper, we have presented two different notions of approximation to the core stability Definitions 1.3 and 3.1. Fain et al. [2018] have proposed an alternative notion of core stability in the context of public goods, allowing deviating voters to utilize the entire budget. They also developed approximate versions of this concept. Based on the existing reduction relation demonstrated in the proof of Theorem 2.9 (see [Mavrov et al., 2023]), we naturally wonder whether there is an inherent connection between the different approximation definitions of the core.

Open Question 7. Is it possible to explore more relations (e.g. reduction relations) between the different notions of approximate core stability?

Moreover, our auditing results are only for the approximate endowment core. Other than closing the log gap between the positive and hardness results in auditing for core stability, we ask the following question:

Open Question 8. Can we audit core stability with the classic utility relaxation? More concretely, given a committee W, can we (approximately) compute the smallest α efficiently such that W does not lie in the α-core (see Definition 1.3)?

7.5 Fairness with Allocation Constraints

In Mavrov et al. [2023], we extend the restrained core to restrained extended justified representation (restrained EJR). We show that Local **pav** satisfies restrained EJR for Approval elections with matroid constraints. However, matroid constraints seem too limited. We conjecture that the restrained EJR always exists, regardless of what constraints we have. Masařík et al. [2023] have recently shown that the conjecture is correct with downward closed constraints, but downward closed constraints still fail to cover some common constraint scenarios like bundled candidates (c_a and c_b can only be chosen or unchosen together). This begs the following question:

Open Question 9. Does a committee with restrained EJR always exist for arbitrary constraints, especially non-downward closed constraints?

Finally, it will be interesting to study fairness with allocation constraints in other economic contexts, for instance, envy-freeness in allocation of private goods, or stability in matchings. The two-stage approach to define the restrained core may widely extend to establish fairness notions for these problems.

REFERENCES

2023. Participatory Budgeting. https://council.nyc.gov/pb/. https://council.nyc.gov/pb/ Accessed: 2023-09-11.

Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. 2010. Correlation Robust Stochastic Optimization. In *SODA*. 1087–1096.

Kenneth J. Arrow and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265–290.

Haris Aziz, Felix Brandt, Edith Elkind, and Piotr Skowron. 2019. *Computational Social Choice: The First Ten Years and Beyond*. Springer International Publishing, 48–65.

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

Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez Fernández, and Piotr Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *AAAI*. 902–909.

Haris Aziz, Jérôme Lang, and Jérôme Monnot. 2016. Computing Pareto Optimal Committees. In *IJCAI*, Subbarao Kambhampati (Ed.). 60–66.

Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018a. Finding Fair and Efficient Allocations. In *EC*. 557–574.

Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018b. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (Ithaca, NY, USA) *(EC '18)*. Association for Computing Machinery, New York, NY, USA, 557–574.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2009. On concentration of self-bounding functions. (2009).

William C. Brainard and Herbert E. Scarf. 2005. How to Compute Equilibrium Prices in 1891. *American Journal of Economics and Sociology* 64, 1 (2005), 57–83.

Steven J. Brams, D. Marc Kilgour, and M. Remzi Sanver. 2007. A Minimax Procedure for Electing Committees. *Public Choice* 132, 3 (2007), 401–420.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of Computational Social Choice* (1st ed.). Cambridge University Press, USA.

Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. 2017. Phragmén's Voting Methods and Justified Representation. In *AAAI*. 406–413.

Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. 2020. Approval-Based Apportionment. In *AAAI*. 1854–1861.

Markus Brill, Paul Gölz, Dominik Peters, Ulrike Schmidt-Kraepelin, and Kai Wilker. 2024. Approval-based apportionment. *Mathematical Programming* 203, 1 (2024), 77–105.

Markus Brill and Jannik Peters. 2023. Robust and verifiable proportionality axioms for multiwinner voting. *arXiv preprint arXiv:2302.01989* (2023).

Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. 2011. Maximizing a Monotone Submodular Function Subject to a Matroid Constraint. *SIAM J. Comput.* 40, 6 (2011), 1740–1766.

Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Econ. Comput.* 7, 3, Article 12 (sep 2019).

John R. Chamberlin and Paul N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *The American Political Science Review* 77, 3 (1983), 718–733.

Nikhil Chandak, Shashwat Goel, and Dominik Peters. 2024. Proportional aggregation of preferences for sequential decision making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9573–9581.

Moses Charikar. 2000. Greedy Approximation Algorithms for Finding Dense Components in a Graph. In *Approximation Algorithms for Combinatorial Optimization*. 84–95.

Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. 2014. Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes. *SIAM J. Comput.* 43, 6 (2014), 1831–1879.

Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally Fair Clustering. In *ICML*. 1032–1041.

Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Group Fairness in Committee Selection. *ACM Trans. Economics and Comput.* 8, 4 (2020), 23:1–23:18.

Richard Cole and Vasilis Gkatzelis. 2015. Approximating the Nash Social Welfare with Indivisible Items. *SIGecom Exch.* 14, 1 (Nov. 2015), 84–88.

H. R. Droop. 1881. On Methods of Electing Representatives. *Journal of the Statistical Society of London* 44, 2 (1881), 141–202.

Devdatt Dubhashi and Alessandro Panconesi. 2009. *Concentration of Measure for the Analysis of Randomized Algorithms* (1st ed.). Cambridge University Press, USA.

Edmund Eisenberg and David Gale. 1959. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.

Ulle Endriss. 2017. *Trends in Computational Social Choice*. Lulu.com.

Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair Allocation of Indivisible Public Goods. In *EC*.

Uriel Feige. 2006. On Maximizing Welfare When Utility Functions Are Subadditive. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing* (Seattle, WA, USA) *(STOC '06)*. 41–50.

Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. 1996. Interactive Proofs and the Hardness of Approximating Cliques. *J. ACM* 43, 2 (mar 1996), 268–292.

Luis Sánchez Fernández, Edith Elkind, Martin Lackner, Norberto Fernández García, Jesús Arias-Fisteus, Pablo Basanta-Val, and Piotr Skowron. 2017. Proportional Justified Representation. In *AAAI*. 670–676.

Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. 2019. Fair Knapsack *(AAAI'19/IAAI'19/EAAI'19)*. AAAI Press, Article 239, 8 pages. https://doi.org/10.1609/aaai.v33i01.33011941

Duncan K Foley. 1970. Lindahl's Solution and the Core of an Economy with Public Goods. *Econometrica: Journal of the Econometric Society* (1970), 66–72.

Eric J. Friedman, Vasilis Gkatzelis, Christos-Alexandros Psomas, and Scott Shenker. 2019. Fair and Efficient Memory Sharing: Confronting Free Riders. In *AAAI*. 1965–1972.

Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately stable committee selection. In *STOC*. 463–472.

Yusuf Kalayci, David Kempe, and Vikram Kher. 2024. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9815–9823.

Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. In *ICML*. 2564–2572.

Mayuresh Kunjir, Brandon Fain, Kamesh Munagala, and Shivnath Babu. 2017. ROBUS: Fair Cache Allocation for Data-parallel Workloads. In *SIGMOD*. 219–234.

Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer.

Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2001. Combinatorial Auctions with Decreasing Marginal Utilities. In *Proceedings of the 3rd ACM Conference on Electronic Commerce* (Tampa, Florida, USA) *(EC '01)*. 18–28.

Erik Lindahl. 1958. Just taxation—a positive solution. In *Classics in the theory of public finance*. 168–176.

Tomáš Masařík, Grzegorz Pierczyński, and Piotr Skowron. 2023. A generalised theory of proportionality in collective decision making. *arXiv preprint arXiv:2307.06077* (2023), 18.

Ivan-Aleksandar Mavrov, Kamesh Munagala, and Yiheng Shen. 2023. Fair Multiwinner Elections with Allocation Constraints. *arXiv preprint arXiv:2305.02868* (2023).

Vahab Mirrokni, Michael Schapira, and Jan Vondrák. 2008. Tight Information-Theoretic Lower Bounds for Welfare Maximization in Combinatorial Auctions. In *Proceedings of the 9th ACM Conference on Electronic Commerce* (Chicago, Il, USA) *(EC '08)*. Association for Computing Machinery, New York, NY, USA, 70–77.

Burt L. Monroe. 1995. Fully Proportional Representation. *The American Political Science Review* 89, 4 (1995), 925–940.

Kamesh Munagala, Yiheng Shen, and Kangning Wang. 2022a. Auditing for core stability in participatory budgeting. In *International Conference on Web and Internet Economics*. Springer, 292–310.

Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022b. Approximate Core for Committee Selection via Multilinear Extension and Market Clearing. In *SODA*. 2229–2252.

Kamesh Munagala, Zeyu Shen, and Kangning Wang. 2021. Optimal Algorithms for Multiwinner Elections and the Chamberlin-Courant Rule. In *EC*. 697–717.

John F. Nash. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), 155–162.

Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional Participatory Budgeting with Additive Utilities. In *NeurIPS*.

Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *EC*. 793–794.

Grzegorz Pierczyński and Piotr Skowron. 2022. Core-stable committees under restricted domains. In *International Conference on Web and Internet Economics*. Springer, 311–329.

Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. 2017. Proportional justified representation. In *AAAI*, Vol. 31.

Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. 2015. Achieving fully proportional representation: Approximability results. *Artificial Intelligence* 222 (2015), 67–103.

Thorvald N Thiele. 1895. Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger* 1895 (1895), 415–441.

Jan Vondrák. 2008. Optimal Approximation for the Submodular Welfare Problem in the Value Oracle Model. In *STOC*. 67–74.

Jan Vondrák. 2010. A note on concentration of submodular functions. *arXiv preprint arXiv:1005.2791* (2010).

Qiqi Yan. 2011. Mechanism Design via Correlation Gap. In *SODA*. 710–719.