Project-Fair and Truthful Mechanisms for Budget Aggregation

Rupert Freeman[1], **Ulrike Schmidt-Kraepelin**[2]

[1]University of Virginia, Charlottesville, VA, USA
[2]TU Eindhoven, The Netherlands
freemanr@darden.virginia.edu, u.schmidt.kraepelin@tue.nl

Abstract

We study the budget aggregation problem in which a set of strategic voters must split a finite divisible resource (such as money or time) among a set of competing projects. Our goal is twofold: We seek truthful mechanisms that provide fairness guarantees to the projects. For the first objective, we focus on the class of moving phantom mechanisms, which are – to this day – essentially the only known truthful mechanisms in this setting. For project fairness, we consider the mean division as a fair baseline, and bound the maximum difference between the funding received by any project and this baseline. We propose a novel and simple moving phantom mechanism that provides optimal project fairness guarantees. As a corollary of our results, we show that our new mechanism minimizes the ℓ_1 distance to the mean for three projects and gives the first non-trivial bounds on this quantity for more than three projects.

1 Introduction

In the budget aggregation problem, a fixed amount of a divisible resource (such as money or time) must be allocated among m competing projects based on the divisions proposed by a set of n voters. For example, in participatory budgeting (Cabannes 2004; Aziz and Shah 2021), citizens vote directly on how a public budget should be divided between a set of public projects. Other examples might include a university department allocating discretionary funding among different initiatives or a group of conference organizers deciding how to divide time among activities such as talks, posters, and social events.

A common and natural solution to this problem is to divide the resource according to the (arithmetic) mean of the votes, guaranteeing that the funding received by each project is proportional to the total support that project receives from the voters. However, using the mean as a budget aggregation rule is not strategyproof.[1] For example, voters can overstate their preference for their favorite projects to bring the funding for that project towards the voter's true preference.

In pursuit of strategyproof mechanisms, Freeman et al. (2021) defined the class of *moving phantom mechanisms*, a high-dimensional generalization of the well-known class of generalized median mechanisms for strategyproof aggregation in one dimension (Moulin 1980). Moving phantom mechanisms are strategyproof when voters have disutilities given by the ℓ_1 distance between their vote and the aggregate division. One particularly natural mechanism, which turns out to be a member of this class, is the one that minimizes the sum of disutilities of the voters (Lindner, Nehring, and Puppe 2008; Goel et al. 2019). Although this rule can be effective, it can also produce outcomes that differ significantly from the mean (for intuition, consider the median in one dimension). If the mean is considered a desirable outcome, then it would be beneficial to discover strategyproof mechanisms that are more aligned with it.

Freeman et al. (2021) introduced the *Independent Markets* mechanism, which is guaranteed to agree with the mean when all voters want to fund only a single project (a mechanism with this property is said to be *proportional*). However, as Caragiannis, Christodoulou, and Protopapas (2022) showed in subsequent work, on other inputs it may produce outcomes that are far from the mean according to the ℓ_1 distance. They propose a different moving phantom mechanism, the *Piecewise Uniform mechanism*, which never outputs budget divisions that have an ℓ_1 distance from the mean larger than $\frac{2}{3}$ when there are only three projects.[2] No positive results are known for higher numbers of projects.

Caragiannis, Christodoulou, and Protopapas (2022) measure the quality of an outcome by its ℓ_1 distance to the mean. However, this measure does not capture how the deviation from the mean is distributed over the projects. For instance, suppose that the mean division over four projects is given by $(70\%, 10\%, 10\%, 10\%)$ and consider two potential aggregate divisions: $a = (50\%, 30\%, 10\%, 10\%)$ and $b = (60\%, 20\%, 0\%, 20\%)$. For both aggregates, the ℓ_1 distance to the mean is 40%. However, while in division a, the first project is being underfunded relative to the mean by 20% of the budget, in division b, no project is over or underfunded by more than 10% of the budget. In this paper, we complement Caragiannis, Christodoulou, and Protopapas's approach by studying the ℓ_∞ distance to the mean, which can be interpreted as a measure of fairness between

[1]We use the terms "strategyproof" and "truthful" interchangeably.

[2]The proof relies on solving a nonlinear program, so this bound is subject to a small error.

projects. Taking the mean to be a project's "entitlement," by how much does the allocation of any project exceed or fall short of this value?

Given that the projects themselves (or the entities behind them) are typically stakeholders in budget aggregation systems, project-fairness guarantees are important to maintain the confidence of the projects in the system.

Our Contributions. We introduce the notion of project fairness for the budget aggregation problem. While our definition is technically similar to the proportionality measure of Caragiannis, Christodoulou, and Protopapas (2022) in that we are interested in the worst case distance (according to some metric) from the mean, the two metrics can differ substantially in which outcomes they prefer. That said, they are related in that an upper bound on the ℓ_∞ distance implies an upper bound on the ℓ_1 distance and vice versa; see Section 6.

We focus on project fairness for the class of moving phantom mechanisms. Whether there exist (anonymous, neutral, and continuous) strategyproof mechanisms outside of this class remains an intriguing open question. As our main result, we define the *Ladder mechanism*, a new moving phantom mechanism that is guaranteed to output a budget division with ℓ_∞ distance from the mean equal to at most $\frac{1}{2} - \frac{1}{2m}$. This bound is tight for moving phantom mechanisms. We additionally show that, while our mechanism may underfund a project by this amount relative to the mean, it never overfunds a project by more than $1/4$, a property that we show to be common to all proportional mechanisms.

As a corollary of our result, we show that our new mechanism guarantees an ℓ_1 distance from the mean of no more than $\frac{2}{3}$ for instances with three projects, which matches the known lower bound. This closes a (very small) gap that was left open by Caragiannis, Christodoulou, and Protopapas (2022), who obtained an upper bound of $\frac{2}{3} + 10^{-5}$ by a complex proof that involved characterizing worst case instances and then solving a non-linear program. In contrast, our proof is combinatorial and relatively simple in comparison. We additionally obtain non-trivial bounds on the ℓ_1 distance from the mean for 4, 5, and 6 projects. Prior to our work, no mechanisms were known to guarantee an ℓ_1 distance less than a trivial upper bound for more than 3 projects.

Related Work. Portioning, also known as (unbounded) divisible participatory budgeting, is an umbrella term for problems in which a continuous divisible resource must be divided among alternatives. The budget aggregation problem is an example of portioning where voters submit complete budget allocation proposals; in addition to the papers discussed above, Elkind, Suksompong, and Teh (2023) perform an axiomatic analysis of various rules in this setting, and find that the mean performs well relative to the other rules they consider. In particular, it is the only one of the considered rules to satisfy the score representation axiom, a natural proportionality property. Goyal et al. (2023) study mechanisms with low sample complexity in terms of their social welfare approximation guarantees. Other variants of portioning include voters submitting ordinal preferences (Airiau et al. 2023), dichotomous preferences (e.g., Bogomolnaia, Moulin, and Stong 2005; Brandl et al. 2021; Michorzewski, Peters, and Skowron 2020), or more general cardinal utility functions over alternatives (Fain, Goel, and Munagala 2016; Wagner and Meir 2023). For an overview of other models and additional related work in participatory budgeting, we refer to the survey of Aziz and Shah (2021).

For the special case of two projects,[3] moving phantom mechanisms reduce to the generalized median mechanisms of Moulin (1980), which take the median of $n + 1$ fixed "phantom" votes and the n submitted votes. These mechanisms have been extensively studied, most notably in the context of strategyproof facility location (Procaccia and Tennenholtz 2013; Aziz et al. 2021). Connections between generalized median mechanisms and mean approximation in one dimension have also been made previously in various contexts (Renault and Trannoy 2005, 2011; Caragiannis, Procaccia, and Shah 2016; Jennings et al. 2023; Caragiannis, Christodoulou, and Protopapas 2022). All of these papers identify the uniform phantom mechanism, which places phantom votes at uniform intervals of $1/n$, as the most desirable generalized median from this perspective. As with other proportional moving phantom mechanisms in the literature, the Ladder mechanism draws heavy inspiration thereof.

Alternative multidimensional aggregation settings that do not require votes and outcomes to sum to one exist in the literature (e.g., Barberà, Gul, and Stacchetti 1993; Barberà, Massó, and Neme 1997; Border and Jordan 1983; Peters, van der Stel, and Storcken 1992). Typically, strategyproof mechanisms in these models can be decomposed into one-dimensional mechanisms taking a generalized median in every coordinate, which would violate our normalization requirement. Accordingly, these problems are very different to ours from a technical perspective.

2 Preliminaries

For any $k \in \mathbb{N}$, let $[k] = \{1, \dots, k\}$ and $[k]_0 = \{0, 1, \dots, k\}$. We denote by $N = [n]$ the set of voters and by $M = [m]$ the set of projects. For any $m \in \mathbb{N}$, we define $\Delta^{(m)} = \{q \in [0,1]^m \mid \sum_{j \in [m]} q_j = 1\}$ to be the standard simplex. For a set of projects M, each voter indicates their *ideal* budget distribution over the projects, i.e., an element of $\Delta^{(m)}$. Formally, these preferences are summarized in a preference profile P, which is a matrix $P \in [0,1]^{n \times m}$ with $(P_{ij})_{j \in [m]} \in \Delta^{(m)}$ for every $i \in N$. A budget aggregation mechanism $\mathcal{A}$ takes as input a preference profile P and outputs an element from $\Delta^{(m)}$. For a given preference profile P, and for any $j \in [m]$, let $\overline{P}_j = \frac{1}{n} \sum_{i \in [n]} P_{ij}$ be the average support of this project.

Moving Phantom Mechanisms For $n \in \mathbb{N}$, a *phantom system* $\mathcal{F}_n = \{f_k : k \in [n]_0\}$ is a family of functions, where $f_k : [0,1] \to [0,1]$ is a continuous, non-decreasing function with $f_k(0) = 0$ and $f_k(1) = 1$ for each k, and $f_0(t) \geq f_1(t) \geq \cdots \geq f_n(t)$ for all $t \in [0,1]$. Then, for any preference profile P, let $t^* \in [0,1]$ be chosen such that

$$\sum_{j \in [m]} \mathrm{med}(f_0(t^*), \dots, f_n(t^*), P_{1j}, \dots, P_{nj}) = 1,$$

[3]Since we have a normalization constraint, the two-project case has only one degree of freedom.

where "med" is the median. Then, we define

$$\mathcal{A}^{\mathcal{F}_n}(P)_j = \mathrm{med}(f_0(t^*), \ldots, f_n(t^*), P_{1j}, \ldots, P_{nj})$$

and say that $\mathcal{A}^{\mathcal{F}_n}$ *reaches normalization at* t^*. While t^* is not always unique, the resulting budget allocation is unique.

Since phantom systems are defined for fixed $n \in \mathbb{N}$, we are interested in *families* of phantom systems, $\mathcal{F} = \{\mathcal{F}_n \mid n \in \mathbb{N}\}$, and define the *moving phantom mechanism* $\mathcal{A}^{\mathcal{F}}$ by applying mechanism $\mathcal{A}^{\mathcal{F}_n}$ to any profile with n voters.

Following Freeman et al. (2021), we will pictorially represent (snapshots of) moving phantom mechanisms in the following way (see, for example, Figure 1). Projects are represented by vertical bars, with voter reports indicated by black horizontal line segments. The vertical position of the segment indicates the report P_{ij}. Phantom positions are indicated by solid blue lines. On every project, the median of the voter and phantom positions is indicated by a rectangle.

Strategyproofness For any $q \in \Delta^{(m)}$, the disutility of a voter i is assumed to be the ℓ_1 distance from q to its ideal point, i.e., $\sum_{j\in[m]} |P_{ij} - q_j|$. It is known that all moving phantom mechanisms are strategyproof in the sense that no voter can decrease the ℓ_1 distance from the aggregate to their ideal distribution by reporting a distribution that is not their ideal one (Freeman et al. 2021). Note that strategyproofness of moving phantom mechanisms rests crucially on the assumption of ℓ_1 (dis)utilities. We refer the reader to Nehring and Puppe (2019) and Goel et al. (2019) for natural interpretations of this utility model in the budgeting setting (and to Varloot and Laraki (2022) for a setting where a utility model other than ℓ_1 is more appropriate).

Proportionality We say that a voter $i \in N$ is *single minded* if $P_{ij} \in \{0, 1\}$ for all $j \in M$. Freeman et al. (2021) define a budget aggregation mechanism to be *proportional* if, for any profile consisting of single-minded voters only, it holds that $\mathcal{A}(P)_j = \overline{P}_j$ for all $j \in M$.

We now introduce the main novel concept of the paper, i.e., a guarantee for the maximum deviation of the funding received by any project from the funding given to this project by the mean aggregation function (which is not strategyproof). Since these bounds can be made more precise by parameterizing them by n and m, we express the resulting bounds in terms of a function $\alpha(n, m)$.

Project Fairness For a function $\alpha : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$, we say that a budget aggregation mechanism is α-project fair if, for any preference profile P on n voters and m projects, and any $j \in [m]$, it holds that $|\mathcal{A}(P)_j - \overline{P}_j| \le \alpha(n, m)$. In addition, we say that a mechanism overfunds by at most α if $\mathcal{A}(P)_j - \overline{P}_j \le \alpha(n, m)$ for all $j \in [m]$, and it underfunds by at most α if $\overline{P}_j - \mathcal{A}(P)_j \le \alpha(n, m)$ for all $j \in [m]$. Clearly, a mechanism is α-project fair if and only if it overfunds by at most α and underfunds by at most α. For simplicity, any function α in this paper maps from $\mathbb{N} \times \mathbb{N}$ to $\mathbb{R}$.

3 Lower Bounds

In this section, we provide several lower bounds on the α-project fairness for (subclasses of) moving phantom mechanisms. This paves the way to the introduction of a new mechanism guaranteeing optimal project fairness. We say that a budget aggregation mechanism is *zero unanimous* if it never funds a project that every voter agrees should receive zero funding, i.e., $P_{ij} = 0$ for all $i \in [n]$ implies that $\mathcal{A}(P)_j = 0$. Zero unanimity is a restriction of the score unanimity condition of Elkind, Suksompong, and Teh (2023), which says that whenever all agents unanimously agree on the funding for a particular project, then this project should receive exactly that level of funding. In Proposition 1, we start by providing a lower bound on the overfunding guarantee of any zero-unanimous moving phantom mechanism.[4]

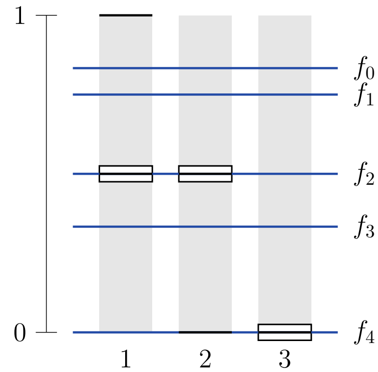

Figure 1: Example from Proposition 1 for $n = 4$ and $m = 3$. See Section 2 for an explanation of how to read our figures. Note that the black line segments each represent two voters who both make the same report.

Proposition 1. *Let $\mathcal{A}^{\mathcal{F}}$ be a zero-unanimous moving phantom mechanism. Then, there exists no α satisfying*

- $\alpha(n, m) < \frac{1}{4}$ *for any $n, m \in \mathbb{N}$, where n is even, or*
- $\alpha(n, m) < \frac{1}{4}\left(1 - \frac{1}{n}\right)$ *for any $n, m \in \mathbb{N}$, where n is odd,*

such that $\mathcal{A}^{\mathcal{F}}$ overfunds by at most α.

Proof. For any $n, m \in \mathbb{N}$, consider the instance in which the voters in $N_1 = [\lfloor n/2 \rfloor]$ report $P_{i1} = 1$ and $P_{ij} = 0$ for all $j \in M \setminus \{1\}$ and the voters in $N \setminus N_1$ report $P_{i1} = P_{i2} = 1/2$ and $P_{ij} = 0$ $j \in M \setminus \{1, 2\}$. By zero unanimity, $\mathcal{A}^{\mathcal{F}}(P)_j = 0$ for all $j \in M \setminus \{1, 2\}$, and therefore $\mathcal{A}^{\mathcal{F}}$ reaches normalization when $f_{\lfloor n/2 \rfloor}(t) = 1/2$, returning the budget distribution $\mathcal{A}^{\mathcal{F}}(P)_1 = \mathcal{A}^{\mathcal{F}}(P)_2 = 1/2$. We refer to Figure 1 for an illustration of the case $n = 4$ and $m = 3$. Now, when n is even, it holds that $\overline{P}_2 = \frac{1}{4}$ and therefore $\mathcal{A}^{\mathcal{F}}(P)_2 - \overline{P}_2 = 1/4$. When n is odd, we get that $\overline{P}_2 = \frac{1}{4}\left(1 + \frac{1}{n}\right)$ and therefore $\mathcal{A}^{\mathcal{F}}(P)_2 - \overline{P}_2 = \frac{1}{4}\left(1 - \frac{1}{n}\right)$. ☐

We continue by providing a lower bound on the underfunding guarantee that any moving phantom mechanism can provide. To this end, we use an example provided by Caragiannis, Christodoulou, and Protopapas (2022).

[4]Dropping zero unanimity allows for moving phantom mechanisms with better overfunding guarantees. For example, the mechanism that always outputs $\frac{1}{m}$ for every project never overfunds by more than $\frac{1}{m}$. That said, this mechanism clearly suffers from high underfunding and completely ignores the voters' preferences.

Proposition 2. *Let $\mathcal{A}^{\mathcal{F}}$ be a moving phantom mechanism. Then, there exists no α such that*

- $\alpha(n,m) < \frac{1}{2}\left(1 - \frac{1}{m}\right)$ *for any $n, m \in \mathbb{N}$, n even, or*
- $\alpha(n,m) < \frac{1}{2}\left(1 - \frac{1}{m}\right)\left(1 - \frac{1}{n}\right)$ *for $n, m \in \mathbb{N}$, n odd,*

and $\mathcal{A}^{\mathcal{F}}$ underfunds by at most α.

Proof. For any $n, m \in \mathbb{N}$, let $N_1 = [\lfloor n/2 \rfloor]$ and $N_2 = N \setminus N_1$. Then, define the profile

$$P_{ij} = \begin{cases} 1 & \text{if } i \in N_1, j = 1 \\ 0 & \text{if } i \in N_1, j \neq 1 \\ 1/m & \text{if } i \in N_2. \end{cases}$$

Caragiannis, Christodoulou, and Protopapas (2022, Theorem 7) prove that for this profile under the restriction that n is even, any moving phantom mechanism returns $\mathcal{A}^{\mathcal{F}}(P)_j = 1/m$ for all $j \in [m]$. It is easy to verify by going through their arguments that the same holds when n is odd. Hence, we receive the following lower bounds for the underfunding guarantees: For n even, it holds that $\overline{P}_1 = \frac{1}{2} + \frac{1}{2m}$, which implies $\overline{P}_1 - \mathcal{A}^{\mathcal{F}}(P)_1 = \frac{1}{2} - \frac{1}{2m}$. For odd n, we get that $\overline{P}_1 = \frac{1}{2}\big(1 + \frac{1}{m}\big(1 + \frac{1}{n}\big) - \frac{1}{n}\big)$ and therefore $\overline{P}_1 - \mathcal{A}^{\mathcal{F}}(P)_1 = \frac{1}{2}\big(1 - \frac{1}{m}\big(1 - \frac{1}{n}\big) - \frac{1}{n}\big)$. □

In Theorem 7 of the next section, we show that any proportional mechanism overfunds by at most $\frac{1}{4}$. Hence, we can focus on finding a proportional mechanism with optimal underfunding guarantee. While doing so, we first seek to understand the space of mechanisms that are optimal for large m, i.e., mechanisms with an underfunding guarantee α satisfying $\lim_{m\to\infty} \alpha(n,m) = \frac{1}{2}$. In Proposition 3, we exhibit a class of moving phantom mechanisms that do not provide an optimal asymptotic underfunding guarantee. As we show in Corollary 4, this class includes the *Independent Markets mechanism*, which has been previously studied by Freeman et al. (2021) and Caragiannis, Christodoulou, and Protopapas (2022). Intuitively, Proposition 3 implies that any mechanism that moves a phantom with high index while the symmetric phantom of low index is still low has to have a higher asymptotic underfunding guarantee than $\frac{1}{2}$.

Proposition 3. *Let $\mathcal{A}^{\mathcal{F}}$ be a moving phantom mechanism and $k \in [\lfloor n/2 \rfloor]_0$. Then, for any $t \in [0,1]$ such that $f_{n-k}(t) > 0$, there exists no α such that*

$$\lim_{m\to\infty} \alpha(n,m) \leq \frac{n-k}{n} - f_k(t)$$

and $\mathcal{A}^{\mathcal{F}}$ underfunds by at most α.

Proof. Let $\mathcal{A}^{\mathcal{F}}$, k, and t be as in the proposition assumptions. Now, for any $m \in \mathbb{N}$ such that $f_k(t) + (m-1) \cdot f_{n-k}(t) > 1$ we can construct a simple instance in which $n-k$ voters cast the vote $(1, 0, \ldots, 0)$ and the remaining k voters cast the vote $(0, \frac{1}{m-1}, \ldots, \frac{1}{m-1})$. By construction, the mechanism $\mathcal{A}^{\mathcal{F}}$ is normalized for some $t' < t$. As a result, we get that

$$\overline{P}_1 - \mathcal{A}^{\mathcal{F}}(P)_1 \geq \frac{n-k}{n} - f_k(t),$$

and the proposition statement follows. □

Below, we show the implication for the *Independent Markets* mechanism, which is defined by the phantom system[5]

$$\mathcal{F}_n = \left\{ f_k(t) = t \cdot \frac{n-k}{n} \text{ for all } k \in [n]_0, t \in [0,1] \right\},$$

for all $n \in \mathbb{N}$.

Corollary 4. *For the Independent Markets mechanism and any $\epsilon > 0$, there exists no function α satisfying*

$$\lim_{m\to\infty} \alpha(n,m) \leq (1-\epsilon)\frac{n-1}{n}$$

such that the Independent Markets mechanism underfunds by at most α.

Proof. For any $\epsilon > 0$ it holds that $f_1(\epsilon) = \epsilon\frac{n-1}{n}$ and $f_{n-1}(\epsilon) = \epsilon\frac{1}{n} > 0$. Therefore, Proposition 3 implies that there exists no α with $\lim_{m\to\infty} \alpha(n,m) \leq \frac{n-1}{n} - \epsilon\frac{n-1}{n}$, such that the Independent Markets mechanism underfunds by at most α. □

4 The Ladder Mechanism

Proposition 3 narrows down the space of moving phantom mechanisms that can achieve a project fairness guarantee of $\frac{1}{2}$ in the limit as m grows: At any moment in time $t \in [0,1]$ when $f_{n-k}(t) > 0$ for any $k \in [\lfloor n/2 \rfloor]_0$, it needs to hold that $f_k(t) \geq \frac{1}{2} - \frac{k}{n}$. Thus, we aim to construct mechanisms that first move the upper phantoms while keeping the lower phantoms at zero. However, while doing so we have to be careful. For example, it might be tempting to consider the moving phantom mechanism which starts by increasing f_0 from 0 to 1 (while keeping all other phantoms at 0), then moves f_1 from 0 to $\frac{n-1}{n}$, and so on. However, the large gap between the middle phantoms leads to problems itself: For any odd n, there exists a profile[6] with $m = 3$ in which this algorithm underfunds a project by $\frac{1}{2} - \frac{1}{2n}$, which is larger than the lower bound from Proposition 2.

There exists one moving phantom mechanism in the literature that avoids both of the described issues: Caragiannis, Christodoulou, and Protopapas (2022) introduced the *Piecewise Uniform mechanism*, which in a first phase spreads the upper $\lceil \frac{n+1}{2} \rceil$ phantoms uniformly within the interval $[0,1]$, and, in a second phase, spreads the lower $\lfloor \frac{n+1}{2} \rfloor$ phantoms uniformly within the interval $[0, \frac{1}{2}]$ while pushing the first half of the phantoms into the interval $[\frac{1}{2}, 1]$. The Piecewise Uniform mechanism avoids the issue captured by Proposition 3 and at no time creates too large a gap between consecutive phantoms. Hence, the mechanism is in fact a promising

[5]This does not exactly fit the definition of a moving phantom mechanism since $f_k(1) < 1$ for $k \in [n]$. However, it is known that $f_k(1) \geq 1 - \frac{k}{n}$ is sufficient to always achieve normalization (Freeman et al. 2021), thus moving all phantoms to 1 is redundant.

[6]Let $n \in \mathbb{N}$ be odd, $m = 3$. Then, $\lfloor n/2 \rfloor$ of the voters report $(1,0,0)$ while $\lceil n/2 \rceil$ of the voters report $(0, \frac{1}{2}, \frac{1}{2})$. The mechanism described above would output $(0, \frac{1}{2}, \frac{1}{2})$ and therefore underfund project 1 by $\frac{1}{2} - \frac{1}{2n}$.The same profile provides a counter example for the mechanism maximizing utilitarian social welfare (Freeman et al. 2021).

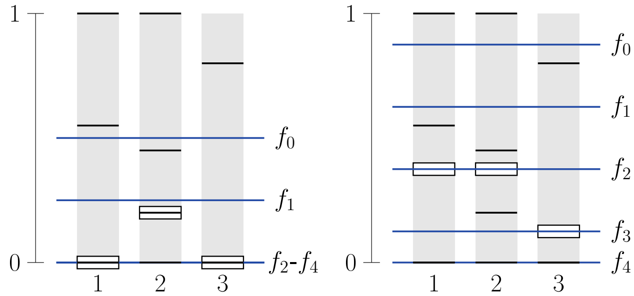

Figure 2: Example execution of the Ladder mechanism with $n = 4$ voters and $m = 3$ projects. The left panel shows the positions of the phantoms at $t = \frac{1}{2}$ (before normalization is reached) while the right panel shows them at $t = \frac{11}{12}$ (exactly when normalization is reached).

candidate for optimal project fairness. That said, the precise definition of the mechanism is intricate, making it difficult to analyze. Instead, we propose a novel and arguably simpler mechanism which also avoids both issues and additionally allows for an elegant proof of optimal project fairness.

We refer to our mechanism as the *Ladder mechanism*, and there are two ways to gain intuition for it: The first view, giving the mechanism its name, thinks of a rope ladder where the ladder rungs correspond to the phantoms. The ladder is then pulled up by its top rung. The second view, being closer to its formal definition, imagines the phantoms being uniformly spread within the interval $[-1, 0]$, and then, as t increases, being pushed upwards (with equal speed) until they are uniformly spread in $[0, 1]$. However, since phantoms need to be non-negative, they only become "active" once they cross 0, which is ensured by the max function in the following definition.

Definition 5 (Ladder Mechanism)**.** *The Ladder mechanism is the moving phantom mechanism defined by the following phantom system for any* $n \in \mathbb{N}$*:*

$$f_k(t) = \max\left(t - \frac{k}{n}, 0\right) \text{ for all } k \in [n]_0, t \in [0, 1].$$

We illustrate the Ladder mechanism in Figure 2. The example displayed in the figure has four voters with reports $(0, 0.2, 0.8), (1, 0, 0), (0, 1, 0), (0.55, 0.45, 0)$. Normalization is reached at $t = \frac{11}{12}$, returning the budget distribution $(\frac{5}{12}, \frac{5}{12}, \frac{1}{6})$.

5 Upper Bounds

In this section, we present our main results, i.e., upper bounds for the overfunding and underfunding guarantees provided by the Ladder mechanism that are *essentially* tight. In this context, we write *essentially* in order to refer to the fact that there is a small gap of $\mathcal{O}(\frac{1}{n})$ between the upper and lower bounds only in the case when n is odd.

To prove the overfunding guarantee (Theorem 7), we need the fact that any moving phantom mechanism is monotone. That is, if a single voter increases its report for a single project while it decreases its report for all other projects, then this project receives at least as much funding as in the original instance. Formally, a budget aggregation mechanism is *monotone* if, for any two profiles P and P' for which there exists a voter i_0 and a project j_0 such that $(P_{ij})_{j\in[m]} = (P'_{ij})_{j\in[m]}$ for all $i \in [n] \setminus i_0$, and $P_{i_0,j_0} < P'_{i_0,j_0}$ while $P_{i_0,j} \geq P'_{i_0,j}$ for all $j \in [m] \setminus j_0$, it holds that $\mathcal{A}^\mathcal{F}(P)_{j_0} \leq \mathcal{A}^\mathcal{F}(P')_{j_0}$.

Lemma 6. *Any moving phantom mechanism is monotone.*

Proof. Let $t \in [0, 1]$ ($t' \in [0, 1]$, respectively) be the time at which mechanism $\mathcal{A}^\mathcal{F}$ reaches normalization on profile P (P', respectively). If $t \leq t'$, then $\mathcal{A}^\mathcal{F}(P)_{j_0} \leq \mathcal{A}^\mathcal{F}(P')_{j_0}$ since phantoms and voters on project j_0 are all weakly higher for P' than for P at the time of normalization. If $t > t'$, then voters and phantoms are weakly lower for P' than for P for all $j \in [m] \setminus \{j_0\}$ at the time of normalization, implying $\mathcal{A}^\mathcal{F}(P)_j \geq \mathcal{A}^\mathcal{F}(P')_j$. By normalization, this implies $\mathcal{A}^\mathcal{F}(P)_{j_0} \leq \mathcal{A}^\mathcal{F}(P')_{j_0}$. □

We are now ready to prove Theorem 7.

Theorem 7. *Let $\mathcal{A}^\mathcal{F}$ be a proportional moving phantom mechanism. Then, $\mathcal{A}^\mathcal{F}$ overfunds by at most α, where*

$$\alpha(n, m) = \begin{cases} \frac{1}{4} & \text{for } n, m \in \mathbb{N}, n \text{ even} \\ \frac{1}{4}\left(1 - \frac{1}{n^2}\right) & \text{for } n, m \in \mathbb{N}, n \text{ odd.} \end{cases}$$

Proof. Consider some profile P with normalization achieved at time t^*, and some project $j \in [m]$. Denote by N^- the set of voters with $P_{ij} < \mathcal{A}^\mathcal{F}(P)_j$ and let $n^- = |N^-|$. Note that for every voter $i \in N^-$ there must exist a project j_i with $P_{ij_i} > \mathcal{A}^\mathcal{F}(P)_{j_i}$, since votes and outputs are normalized. Starting from P, construct a profile P' by, for every voter $i \in N^-$, changing i's vote to be single minded on project j_i. Note that, holding the position of the phantoms fixed at $\{f_k(t^*) : k \in [n]_0\}$, the median on every coordinate is (weakly) lower in P' than in P, with the median on project j being the same in the two profiles. So it might be the case that to achieve normalization in profile P', we need to advance the phantoms to $\{f_k(t') : k \in [n]_0\}$ for some $t' > t^*$. Therefore, $\mathcal{A}^\mathcal{F}(P)_j \leq \mathcal{A}^\mathcal{F}(P')_j$. Let us now construct a profile P'' by starting with P' and, for every voter $i \notin N^-$, setting their vote to be single-minded on project j. By monotonicity, $\mathcal{A}^\mathcal{F}(P'')_j \geq \mathcal{A}^\mathcal{F}(P')_j$. By proportionality of $\mathcal{A}^\mathcal{F}$, we have $\mathcal{A}^\mathcal{F}(P'')_j = 1 - \frac{n^-}{n}$. Combining the inequalities, we get $\mathcal{A}^\mathcal{F}(P)_j \leq \mathcal{A}^\mathcal{F}(P')_j \leq \mathcal{A}^\mathcal{F}(P'')_j = 1 - \frac{n^-}{n}$.

To complete the proof, note that $n - n^-$ voters have report $P_{ij} \geq \mathcal{A}^\mathcal{F}(P)_j$, by the definition of N^-. Therefore, $\overline{P}_j \geq (1 - \frac{n^-}{n})\mathcal{A}^\mathcal{F}(P)_j$. We have

$$\begin{aligned} \mathcal{A}^\mathcal{F}(P)_j - \overline{P}_j &\leq \mathcal{A}^\mathcal{F}(P)_j - \left(1 - \frac{n^-}{n}\right)\mathcal{A}^\mathcal{F}(P)_j \\ &= \frac{n^-}{n}\mathcal{A}^\mathcal{F}(P)_j \leq \frac{n^-}{n}\left(1 - \frac{n^-}{n}\right), \end{aligned}$$

which is at most $\frac{1}{4}$ when n is even and at most $\frac{1}{4}\left(1 - \frac{1}{n^2}\right)$ when n is odd. □

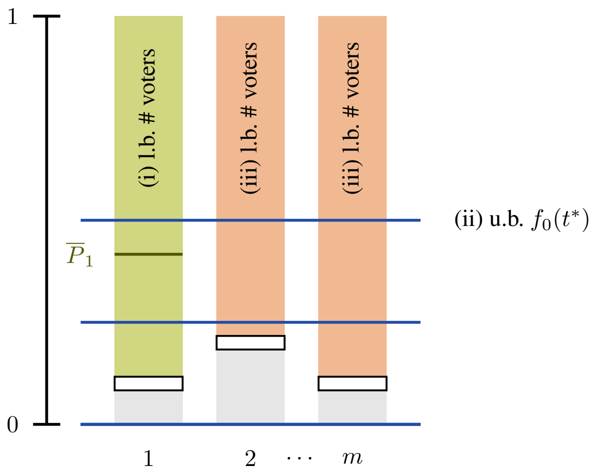

Figure 3: Proof sketch of Theorem 8: We assume for contradiction that the Ladder mechanism underfunds project 1 by more than $\frac{1}{2}(1 - \frac{1}{m})$. We divide the proof into four steps: Using that the mean for project 1 is high, in step (i), we derive a lower bound for the number of voters reporting a value in the interval $(\mathcal{A}^{\mathcal{F}}(P)_1, 1]$ (indicated by green). (ii) Since we know that the total number of phantoms and voters strictly above $\mathcal{A}^{\mathcal{F}}(P)_1$ is at most n, we derive an upper bound for the highest phantom at the point of normalization, i.e., $f_0(t^*)$. (iii) Building upon (ii), we can upper bound the number of phantoms within each interval $[\mathcal{A}^{\mathcal{F}}(P)_j, 1]$ (indicated by orange) and thereby lower bound the number of voters reporting a value in the same interval. This in turn allows us to derive a lower bound on the mean of each project $j \neq 1$. (iv) Summing over all lower bounds on the mean implies a contradiction to the fact that the means sum up to 1.

We can easily verify that the Ladder mechanism satisfies proportionality: Freeman et al. (2021, Section 5) argue that a moving phantom mechanism satisfies proportionality if there exists $t \in [0,1]$ such that $f_k(t) = 1 - \frac{k}{n}$ holds for all $k \in [n]_0$, which is the case for the Ladder mechanism when $t = 1$. Hence, as an immediate corollary of Theorem 7, we get that the overfunding guarantee of the Ladder mechanism is essentially optimal. We now turn to proving our main result, i.e., an essentially tight upper bound for the underfunding guarantee of the Ladder mechanism. We provide a proof sketch in Figure 3.

Theorem 8. *The Ladder mechanism underfunds by at most α, where*

$$\alpha(n,m) = \frac{1}{2}\Big(1 - \frac{1}{m}\Big) \quad \textit{for all } n, m \in \mathbb{N}.$$

Proof. Let $\mathcal{A}^{\mathcal{F}}$ be the Ladder mechanism and let P be a preference profile. For the sake of contradiction, assume that there exists a project $j \in M$ with

$$\overline{P}_j - \mathcal{A}^{\mathcal{F}}(P)_j > \frac{1}{2} - \frac{1}{2m}. \tag{1}$$

We assume without loss of generality that $j = 1$.

We introduce the following notation: For simplicity, we write $a_j = \mathcal{A}^{\mathcal{F}}(P)_j$ for all $j \in M$. For a given project $j \in M$ and some interval $I \subseteq [0,1]$, we denote by $n_j(I)$ the number of agents within the interval I, i.e., $|\{i \in N \mid P_{ij} \in I\}|$. Similarly, we denote by $p(I)$ the number of phantoms in interval I, i.e., $|\{k \in [n]_0 \mid f_k(t^*) \in I\}|$, where $t^* \in [0,1]$ is some arbitrary point of normalization.

Step (i) We start by showing that

$$n_1((a_1,1]) \geq n \cdot \frac{\overline{P}_1 - a_1}{1 - a_1}. \tag{2}$$

This is because, given $n_1((a_1,1])$ voters with report strictly above a_1, the highest possible mean is attained when all of them report 1 and the remaining voters report a_1. Formally,

$$n_1((a_1,1]) \cdot 1 + (n - n_1((a_1,1])) \cdot a_1 \geq n\overline{P}_1.$$

Rearranging this inequality yields Equation (2).

Step (ii) In this step, we derive an upper bound on the value of the highest phantom, i.e., $f_0(t^*)$. By definition of a_1 as the median on the first project, it holds that $n_1((a_1,1]) + p((a_1,1]) \leq n$, which yields an upper bound for the number of phantoms strictly above a_1. Formally,

$$\begin{aligned} &p((a_1,1]) \leq n - n_1((a_1,1]) \\ &\qquad \overset{eq.\ (2)}{\leq} n \cdot \frac{1-a_1}{1-a_1} - n \cdot \frac{\overline{P}_1 - a_1}{1-a_1} = n \cdot \frac{1-\overline{P}_1}{1-a_1}. \end{aligned}$$

Since any two consecutive phantoms are separated by a distance of at most $\frac{1}{n}$, this yields an upper bound for the position of the highest phantom. Namely, the highest phantom is located at a position no greater than $a_1 + \frac{1-\overline{P}_1}{1-a_1}$.

Step (iii) We next derive lower bounds on the mean for any other project $j \in \{2, \ldots, m\}$:

$$\overline{P}_j \geq \Big(1 - a_1 + a_j - \frac{1-\overline{P}_1}{1-a_1}\Big)a_j. \tag{3}$$

As this bound clearly holds in the case that $a_j = 0$, we can assume in the following case distinction that $a_j > 0$.

Case 1: There is no phantom at a_j. We claim the following upper bound on the number of phantoms weakly above a_j:

$$\begin{aligned} p([a_j,1]) &\leq \Big\lceil n\Big(a_1 + \frac{1-\overline{P}_1}{1-a_1} - a_j\Big)\Big\rceil \\ &\leq n\Big(a_1 + \frac{1-\overline{P}_1}{1-a_1} - a_j\Big) + 1 \end{aligned}$$

This bound holds because of the upper bound on the highest phantom and the fact that any two consecutive phantoms above a_j have distance exactly $\frac{1}{n}$. The ceiling function comes from the fact that a_j and the lowest phantom above a_j may have smaller distance.

Case 2: There is a phantom at a_j. We claim the following upper bound on the number of phantoms weakly above a_j:

$$p([a_j,1]) \leq n\Big(a_1 + \frac{1-\overline{P}_1}{1-a_1} - a_j\Big) + 1.$$

This bound holds because of the upper bound on the highest phantom and the fact that any two consecutive phantoms above a_j have a distance exactly $\frac{1}{n}$. Moreover, since $a_j > 0$, this minimum distance also holds for a_j and the smallest phantom above a_j. The $+1$ comes from the phantom at a_j.

Note that we obtained the same bound in both cases. Now, since $n_j([a_j, 1]) + p([a_j, 1]) \geq n + 1$ (as a_j is the median of $2n + 1$ values), we can use the above bound to derive a lower bound on the number of voters weakly above a_j:

$$n_j([a_j, 1]) \geq n+1-p([a_j, 1]) \geq n\big(1-a_1-\frac{1-\overline{P}_1}{1-a_1}+a_j\big).$$

The lowest mean for project j is attained when $n_j([a_j, 1])$ agents report a_j and the remaining agents report 0. Hence, we obtain the following lower bound on the mean:

$$\overline{P}_j \geq \big(1 - a_1 + a_j - \frac{1-\overline{P}_1}{1-a_1}\big)a_j.$$

Step (iv) We complete the proof by summing over the mean of all projects. Using our bounds from step (iii), the assumption that project 1 is underfunded, and that $\sum_{j=1}^{m} a_j = 1$, we get:

$$\begin{aligned}
\overline{P}_1+\sum_{j=2}^{m}\overline{P}_2 &\overset{eq.(3)}{\geq} \overline{P}_1 + \sum_{j=2}^{m}\big(1-a_1+a_j-\frac{1-\overline{P}_1}{1-a_1}\big)a_j \\
&= \overline{P}_1 + \left(1-a_1-\frac{1-\overline{P}_1}{1-a_1}\right)\sum_{j=2}^{m} a_j + \sum_{j=2}^{m} a_j^2 \\
&= 2\overline{P}_1 - 1 + (1-a_1)^2 + \sum_{j=2}^{m} a_j^2 \\
&\geq 2\overline{P}_1 - 1 + (1-a_1)^2 + (m-1)\left(\frac{1-a_1}{m-1}\right)^2 \\
&\overset{eq.(1)}{>} 2\big(a_1 + \frac{1}{2} - \frac{1}{2m}\big) - 1 + \frac{m}{m-1}(1-a_1)^2 \\
&= 2a_1 - \frac{1}{m} + \frac{m}{m-1}(1-a_1)^2 \\
&= \frac{m-1}{m} - 2(1-a_1) + \frac{m}{m-1}(1-a_1)^2 + 1 \\
&= \left(\sqrt{\frac{m-1}{m}} - \sqrt{\frac{m}{m-1}}(1-a_1)\right)^2 + 1 \geq 1,
\end{aligned}$$

where the unlabeled inequality stems from the fact that $\sum_{j=2}^{m} a_j^2$ is the sum of convex functions, minimized when all values a_j are equal. We obtained a contradiction to the fact that the sum of the means is 1. □

6 Implications for ℓ_1 Distance

Caragiannis, Christodoulou, and Protopapas (2022) study the maximum ℓ_1 distance between the output of any truthful budget aggregation mechanism and the mean. More precisely, consider some budget aggregation mechanism $\mathcal{A}$. Then, for some real number[7] $\alpha \in [0, 2]$, $\mathcal{A}$ is said to be *α-approximate for m*, if

$$\sum_{j\in[m]} |\mathcal{A}(P)_j - \overline{P}_j| \leq \alpha$$

holds for any instance P with m projects. Caragiannis, Christodoulou, and Protopapas (2022) show that no moving phantom mechanism can achieve an approximation better than $1 - \frac{1}{m}$ and that the Piecewise Uniform mechanism is $(\frac{2}{3} + \epsilon)$-approximate for $m = 3$ and some $\epsilon \in [0, 10^{-5}]$. Building upon our results from Section 5, we are able to improve upon this result and show that the Ladder mechanism is $\frac{2}{3}$-approximate for $m = 3$ as well as non-trivial upper bounds for larger m. The bounds follow from the following general result relating overfunding and underfunding to approximation guarantees.

Lemma 9. *If a mechanism $\mathcal{A}$ overfunds by at most β and underfunds by at most γ (for fixed $m \in \mathbb{N}$ and all $n \in \mathbb{N}$), then $\mathcal{A}$ is α-approximate for m, where*

$$\alpha = 2 \cdot \max_{k\in[m-1]} \min\{k\beta, (m-k)\gamma\}.$$

Proof. Consider some profile P. Suppose that $\mathcal{A}(P)_j \geq \overline{P}_j$ for k projects and $\mathcal{A}(P)_j < \overline{P}_j$ for $m - k$ projects. Then,

$$\sum_{j:\mathcal{A}(P)_j\geq\overline{P}_j} |\mathcal{A}(P)_j - \overline{P}_j| \leq k\beta,$$

and analogously,

$$\sum_{j:\mathcal{A}(P)_j<\overline{P}_j} |\mathcal{A}(P)_j - \overline{P}_j| \leq (m-k)\gamma.$$

Moreover, since both $\mathcal{A}(P)$ and $\overline{P}$ are normalized,

$$\sum_{j:\mathcal{A}(P)_j\geq\overline{P}_j} |\mathcal{A}(P)_j - \overline{P}_j| = \sum_{j:\mathcal{A}(P)_j<\overline{P}_j} |\mathcal{A}(P)_j - \overline{P}_j|,$$

which together with the bounds above implies

$$\sum_{j\in[m]} |\mathcal{A}(P)_j - \overline{P}_j| \leq 2 \cdot \min\{k\beta, (m-k)\gamma\}.$$

Taking the maximum over all possible choices of k yields the result. Note that $k = 0$ is impossible, while $k = m$ leads to an ℓ_1 distance to the mean of 0. □

The desired result is derived by applying Theorems 7 and 8 and Lemma 9. We summarize our results in Table 1. Note that directly applying this approach for $m > 6$ gives upper bounds that are larger than the trivial upper bound of 2, which is why Corollary 10 applies only to $3 \leq m \leq 6$.

Corollary 10. *The Ladder mechanism is $\frac{2}{3}$-approximate for $m = 3$, 1-approximate for $m = 4$, $\frac{3}{2}$-approximate for $m = 5$, and $\frac{5}{3}$-approximate for $m = 6$.*

Proof. The ladder mechanism overfunds, for any $n \in \mathbb{N}$ and $m \in \mathbb{N}$ by at most $\beta = \frac{1}{4}$ and underfunds for any $n \in \mathbb{N}$ and $m \in \mathbb{N}$ by at most $\gamma = \frac{1}{2} - \frac{1}{2m}$. For $m = 3$, we have

$$2\min\{k\beta, (m-k)\gamma\} = 2\min\{\tfrac{1}{4}, \tfrac{2}{3}\} = \tfrac{1}{2} \text{ for } k = 1$$

[7] In this section, we focus on upper bounds that hold for specific values of m and are independent of n, hence, for the sake of presentation we omit the parameterization of the upper bounds here.

m	Lo. Bound	Up. Bound	Previous Up. Bound
3	$2/3$	$2/3$	$2/3+\epsilon$
4	$3/4$	1	2
5	$4/5$	$3/2$	2
6	$5/6$	$5/3$	2

Table 1: A summary of our results for the worst case ℓ_1 distance from the mean. Lower bounds (holding for any moving phantom mechanism) and the previous $m = 3$ upper bound are from the work of Caragiannis, Christodoulou, and Protopapas (2022). Other previous upper bounds are trivial. In the previous $m = 3$ upper bound, ϵ is some small constant no larger than 10^{-5}.

$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{2},\tfrac{1}{3}\} = \tfrac{2}{3} \text{ for } k=2,$$

hence the ladder mechanism is $\frac{2}{3}$-approximate.

For $m = 4$, we have

$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{4},\tfrac{9}{8}\} = \tfrac{1}{2} \text{ for } k=1$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{2},\tfrac{3}{4}\} = 1 \text{ for } k=2$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{3}{4},\tfrac{3}{8}\} = \tfrac{3}{4} \text{ for } k=3,$$

hence the ladder mechanism is 1-approximate.

For $m = 5$, we have

$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{4},\tfrac{16}{10}\} = \tfrac{1}{2} \text{ for } k=1$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{2},\tfrac{12}{10}\} = 1 \text{ for } k=2$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{3}{4},\tfrac{8}{10}\} = \tfrac{3}{2} \text{ for } k=3$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{1,\tfrac{4}{10}\} = \tfrac{4}{5} \text{ for } k=4,$$

hence the ladder mechanism is $\frac{3}{2}$-approximate.

For $m = 6$, we have

$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{4},\tfrac{25}{12}\} = \tfrac{1}{2} \text{ for } k=1$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{1}{2},\tfrac{20}{12}\} = 1 \text{ for } k=2$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{3}{4},\tfrac{15}{12}\} = \tfrac{3}{2} \text{ for } k=3$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{1,\tfrac{10}{12}\} = \tfrac{5}{3} \text{ for } k=4$$
$$2\min\{k\beta,(m-k)\gamma\} = 2\min\{\tfrac{5}{4},\tfrac{5}{12}\} = \tfrac{5}{6} \text{ for } k=5,$$

hence the ladder mechanism is $\frac{5}{3}$-approximate. □

7 Discussion

We introduce the notion of project fairness for the budget aggregation problem, defined by the maximum difference between the funding that a project receives and the funding that it would have received under the mean division of the budget. Our main technical contribution is to define the Ladder mechanism and show that it achieves essentially tight project fairness bounds. Additionally, our result yields a guarantee on the maximum ℓ_1 distance between the output of the Ladder mechanism and the mean division, which is optimal for $m = 3$ and the first non-trivial guarantees for $m \in \{4, 5, 6\}$.

Several open questions remain. Perhaps most intriguing is whether we can achieve better project-fairness guarantees with strategyproof mechanisms that are not moving phantom mechanisms (a project fairness lower bound of $\alpha(n, m) = \frac{1}{4}$ follows directly from an argument of Caragiannis, Christodoulou, and Protopapas (2022, Theorem 6), but this is far from our upper bounds). Of course, resolving this question in the affirmative would require a resolution to the question of whether moving phantom mechanisms comprise the complete space of strategyproof mechanisms in this setting. It would also be interesting to characterize the class of optimal project-fair moving phantom mechanisms. Finally, it may be possible to tighten our analysis in Section 6. We have made use only of the project fairness bounds locally on each project, but perhaps a more global analysis would do better.

Acknowledgments

Ulrike Schmidt-Kraepelin was supported by the *Deutsche Forschungsgemeinschaft* (under grant BR 4744/2-1), the *Centro de Modelamiento Matemático (CMM)* (under grant FB210005, BASAL funds for center of excellence from ANID-Chile) and *ANID-Chile* (grant ACT210005). Moreover, she was supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778 while being in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI) in Berkeley, California, during the Fall 2023 semester.

We thank an anonymous AAAI reviewer for catching a mistake in an earlier version of Proposition 1.

References

Airiau, S.; Aziz, H.; Caragiannis, I.; Kruger, J.; Lang, J.; and Peters, D. 2023. Portioning using ordinal preferences: Fairness and efficiency. *Artificial Intelligence*, 314: 103809.

Aziz, H.; Lam, A.; Lee, B. E.; and Walsh, T. 2021. Strategyproof and Proportionally Fair Facility Location. *arXiv preprint arXiv:2111.01566*.

Aziz, H.; and Shah, N. 2021. Participatory budgeting: Models and approaches. *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, 215–236.

Barberà, S.; Gul, F.; and Stacchetti, E. 1993. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2): 262–289.

Barberà, S.; Massó, J.; and Neme, A. 1997. Voting under constraints. *Journal of Economic Theory*, 76(2): 298–321.

Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2): 165–184.

Border, K. C.; and Jordan, J. S. 1983. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1): 153–170.

Brandl, F.; Brandt, F.; Peters, D.; and Stricker, C. 2021. Distribution Rules Under Dichotomous Preferences: Two Out of Three Ain't Bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 158–179.

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

Caragiannis, I.; Christodoulou, G.; and Protopapas, N. 2022. Truthful aggregation of budget proposals with proportionality guarantees. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 4917–4924.

Caragiannis, I.; Procaccia, A.; and Shah, N. 2016. Truthful Univariate Estimators. In *Proceedings of the 33rd International Conference on Machine Learning*, 127–135.

Elkind, E.; Suksompong, W.; and Teh, N. 2023. Settling the Score: Portioning with Cardinal Preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence*, 621–628.

Fain, B.; Goel, A.; and Munagala, K. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th International Conference on Web and Internet Economics*, 384–399.

Freeman, R.; Pennock, D. M.; Peters, D.; and Vaughan, J. W. 2021. Truthful aggregation of budget proposals. *Journal of Economic Theory*, 193: 105234.

Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2019. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation*, 7(2): 1–27.

Goyal, M.; Sakshuwong, S.; Sarmasarkar, S.; and Goel, A. 2023. Low Sample Complexity Participatory Budgeting. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming*, 70:1–70:20.

Jennings, A. B.; Laraki, R.; Puppe, C.; and Varloot, E. M. 2023. New characterizations of strategy-proofness under single-peakedness. *Mathematical Programming*, 1–32.

Lindner, T.; Nehring, K.; and Puppe, C. 2008. Allocating Public Goods via the Midpoint Rule. In *Proceedings of the 9th International Meeting of the Society for Social Choice and Welfare*.

Michorzewski, M.; Peters, D.; and Skowron, P. 2020. Price of Fairness in Budget Division and Probabilistic Social Choice. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, 2184–2191.

Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice*, 437–455.

Nehring, K.; and Puppe, C. 2019. Resource allocation by frugal majority rule. Technical report, KIT Working Paper Series in Economics.

Peters, H.; van der Stel, H.; and Storcken, T. 1992. Pareto optimality, anonymity, and strategy-proofness in location problems. *International Journal of Game Theory*, 21: 221–235.

Procaccia, A. D.; and Tennenholtz, M. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation*, 1(4): 1–26.

Renault, R.; and Trannoy, A. 2005. Protecting minorities through the average voting rule. *Journal of Public Economic Theory*, 7(2): 169–199.

Renault, R.; and Trannoy, A. 2011. Assessing the extent of strategic manipulation: the average vote example. *SERIEs*, 2: 497–513.

Varloot, E. M.; and Laraki, R. 2022. Level-strategyproof Belief Aggregation Mechanisms. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 335–369.

Wagner, J.; and Meir, R. 2023. Strategy-Proof Budgeting via a VCG-like Mechanism. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory*, 401–418.