

District-Fair Participatory Budgeting

D. Ellis Hershkowitz,¹ Anson Kahng,¹ Dominik Peters,² Ariel D. Procaccia²

¹ Carnegie Mellon University

² Harvard University

dershko@cs.cmu.edu, akahng@cs.cmu.edu, dpeters@seas.harvard.edu, arielpro@seas.harvard.edu

Abstract

Participatory budgeting is a method used by city governments to select public projects to fund based on residents' votes. Many cities use participatory budgeting at a district level. Typically, a budget is divided among districts proportionally to their population, and each district holds an election over local projects and then uses its budget to fund the projects most preferred by its voters. However, district-level participatory budgeting can yield poor social welfare because it does not necessarily fund projects supported across multiple districts. On the other hand, decision making that only takes global social welfare into account can be unfair to districts: A social-welfare-maximizing solution might not fund any of the projects preferred by a district, despite the fact that its constituents pay taxes to the city. Thus, we study how to fairly maximize social welfare in a participatory budgeting setting with a single city-wide election. We propose a notion of fairness that guarantees each district at least as much welfare as it would have received in a district-level election. We show that, although optimizing social welfare subject to this notion of fairness is NP-hard, we can efficiently construct a lottery over welfare-optimal outcomes that is fair in expectation. Moreover, we show that, when we are allowed to slightly relax fairness, we can efficiently compute a fair solution that is welfare-maximizing, but which may slightly overspend the budget.

1 Introduction

Participatory budgeting is a democratic approach to the allocation of public funds. In the participatory budgeting paradigm, city governments fund public projects based on constituents' votes. In contrast to budget committees, which operate behind closed doors, participatory budgeting promises to directly take the voices of the community into account. Since 2014, Paris has allocated more than €100 million per year using constituents' votes. Many other cities around the globe — including Porto Alegre, New York City, Boston, Chicago, San Francisco, Lisbon, Madrid, Seoul, Chengdu, and Toronto — employ participatory budgeting (Cabannes 2004, 2014; Aziz and Shah 2020).

Typically, participatory budgeting is used at a district-level. Each district of the city is allotted a budget proportional to its size. Constituents living in a given district vote on projects such as park, road or school improvements local to the district, using some version of approval voting. Then, the district's

budget is spent according to these votes. For instance, in Paris a participatory budget is split between 20 districts (a.k.a. *arrondissements*), constituents vote and then each district runs a greedy algorithm to maximize the total social welfare — i.e., the total number of votes — of the funded projects.¹

Having separate elections for each district leads to several problems. Foremost, projects that are not local to a single district cannot be accommodated. For this reason, Paris must run an additional election for city-wide projects. However, this splits the available budget for participatory budgeting between district-level and city-wide elections in an ad hoc manner, which is not informed by votes.² Further, people may have interests in multiple districts, such as those who live and work in different districts. For this reason, Paris has to allow residents to choose the district in which they vote. Lastly, a project that only benefits voters at the edge of a district may receive a number of votes that is not proportional to the number of potential beneficiaries.

A simple solution to these problems is a single city-wide election. However, such a voting scheme may result in unfair outcomes. For instance, if votes are aggregated to maximize social welfare (i.e., as is presently done in Paris on the district level) then it is possible that some districts might have none of their preferred projects funded despite deserving a large proportion of the budget. Such outcomes are likely when some districts are much more populous than others, in which case projects local to small districts cannot gather sufficiently many votes. Ideally, we would like a system that balances the tradeoff between social welfare and fairness without an arbitrary, pre-determined split between district-specific and city-wide funding. This motivates our central research question:

How can we maximize social welfare in a way that is fair to all districts?

Intuitively, a solution that is fair to all districts should somehow represent each districts' constituents. One way to formalize this intuition is to stipulate that no district should be able to obtain higher utility by purchasing projects with its proportional share of the budget. In particular, each district

¹More specifically, projects are selected in descending order of vote count until the budget runs out.

²In 2016, this split in Paris was €64.3 million for district elections and €30 million for city-wide elections (Cabannes 2017).

should receive at least as much utility as it would have received had it held a district-level election with its proportional share of the budget. We call this guarantee *district fairness*.³ A district-fair allocation of funds always exists, since an outcome obtained by holding separate district elections is district fair. We aim to find district-fair outcomes that maximize social welfare. Such an outcome will be a Pareto-improvement on the status quo of district-level participatory budgeting, in the sense that each district’s welfare has increased.

Our Results. In our model we think of (utilitarian) social welfare as induced by a given value assigned by each district to each project; our goal is to maximize the sum of these values over districts and selected projects. Note that this model captures the setting of approval votes, where each voter decides on a collection of projects to vote for; the social welfare of a district for a project would then be interpreted as the project’s overall number of approvals from voters in that district. This observation is important because some variant of approval voting is used in most real-world participatory budgeting elections, including in Paris.

We also assume that each district is endowed with an arbitrary fraction of the total budget. Clearly this captures, as a special case, the common setting where the endowment of each district is proportional to its size. Moreover, the reasoning behind the existence of district-fair outcomes immediately applies to the more general setting.

We first show that it is NP-complete to compute an allocation that is welfare-maximizing subject to district fairness. This result holds even for the case of approval votes and proportional budgets, and therefore the generality of our model only strengthens our positive (algorithmic) results without weakening the main negative (hardness) result. We also show that the natural linear program (LP) formulation of the problem has an unbounded integrality gap. Since participatory budgeting elections can be large — hundreds of projects are proposed and hundreds of thousands of votes are cast in Paris — computational complexity can become a problem in practice. Thus, we seek polynomial-time solutions with reasonable approximation guarantees.

There are several ways one might relax our problem or trade-off between parameters in our problem. In this work, we design polynomial-time algorithms that work when we relax or approximate some of the following: (1) the achieved social welfare; (2) the spent budget; (3) the fairness of the solution; and (4) the absence of randomization.

We first relax (4) by considering distributions over outcomes, a.k.a. “lotteries”. We show that using a multiplicative-weights-type algorithm, one can efficiently find a lottery that guarantees budget feasibility (ex post), optimum social welfare (ex ante), and district-fairness in expectation up to an ϵ (ex ante). Since the fairness guarantee only holds in expectation, some districts may be underserved once the lottery is realized. However, since participatory budgeting typically happens repeatedly (e.g., annually), such districts could be

compensated in the next election, for example by increasing their share of the budget in the next year.

We next consider what sort of deterministic guarantees are achievable. To this end, we show how to use techniques from submodular optimization to find an outcome that is district fair “up to one project” and which achieves optimum social welfare with the caveat that the outcome may need to spend 64.7% more money than was originally budgeted. We also give a randomized algorithm with the same guarantees but which overshoots the budget by only a $1/e$ ($\approx 37\%$) fraction with high probability. Additionally, as a corollary of these results, we give both deterministic and randomized algorithms that achieve weaker utility and fairness guarantees but do not overspend the available budget.

Related Work. The social choice literature on participatory budgeting has both studied the voting rules used in practice, and designed original voting schemes. Goel et al. (2019) study knapsack voting, used for example in Madrid (Cabannes 2014), where voters cannot approve more projects than fit into the budget constraint. Talmon and Faliszewski (2019) axiomatically study a variety of approval-based rules that maximize social welfare, both greedy and optimal ones.

The unit cost case (where all projects have the same cost) is best-studied, as *multi-winner* or *committee* elections (Faliszewski et al. 2017). For example, this setting models the election of a parliament. A main focus of that literature is the computational complexity of the winner determination of various voting rules. More relevant for our purposes are fairness axioms used in this setting. The most prominent such axioms are variants of *justified representation* (Aziz et al. 2017). These axioms are formulated for approval votes, and require that arbitrary subgroups of the electorate need to be represented in the outcome if they are *cohesive*, in the sense that there are a sufficient number of projects that are approved by every member of the subgroup. Several voting rules are known to satisfy these conditions, including Phragmén’s rule and Thiele’s Proportional Approval Voting (Janson 2016; Sánchez-Fernández et al. 2017; Brill et al. 2017; Aziz et al. 2018). By contrast, district-fairness gives guarantees to a specific selection of subgroups (i.e., disjoint districts) but does not require these groups to be cohesive.

A very strong fairness axiom that is sometimes discussed in the context of committee elections and participatory budgeting is the *core* (Fain, Goel, and Munagala 2016; Aziz et al. 2017; Fain, Munagala, and Shah 2018). It insists that every subgroup (or *coalition*) must be represented (in the sense that it should not be possible for the subgroup to propose an alternative use of their proportional share of the budget that each group member prefers to the chosen outcome), without a cohesiveness requirement. For approval-based elections, it is a major open question whether there always exists a core outcome. For general additive utilities, there are instance where no core outcome exists (Fain, Munagala, and Shah 2018), but several researchers have proved the existence of approximations to the core (Jiang, Munagala, and Wang 2020; Fain, Munagala, and Shah 2018; Cheng et al. 2019; Peters and Skowron 2020). A district-fair outcome is, in a sense, in the

³Our notion of district fairness can be thought of as a form of *individual rationality* where every district is seen as an “individual.”

core: no subgroup which coincides with a district can block the outcome. Thus, our work shows that for general utilities, a core-like outcome exists if we only allow a specific collection of (disjoint) coalitions to block.

The problem of *knapsack sharing* (Brown 1979) has a similar motivation to our problem. The knapsack sharing problem supposes that the *projects* are separated into districts (instead of, in our case, the voters), and each project comes with a cost and a value. The aim is to find a budget-feasible set of projects that maximize the minimum total value of the projects in a district. Note that in this formulation all districts are treated equally (there is no weighting by district population) and that there is no notion of the value of a project to a specific district. The literature contains a variety of algorithms for solving this NP-hard problem (e.g., Yamada and Futakawa 1997; Yamada, Futakawa, and Kataoka 1998; Hifi, M'Halla, and Sadfi 2005; Fujimoto and Yamada 2006).

2 Formal Problem, Notation and Definitions

Formally, the setting we consider is as follows. We are given a budget $b \in \mathbb{Z}_{\geq 1}$. There are m possible projects $\mathcal{P} = \{x_1, \dots, x_m\}$ with associated nonnegative costs $c : \mathcal{P} \rightarrow \mathbb{Z}_{\geq 0}$. We refer to a subset $W \subseteq \mathcal{P}$ as an *outcome*. The cost of an outcome W is $c(W) := \sum_{x_j \in W} c(x_j)$. We say that a subset W is *budget-feasible* if $c(W) \leq b$.

There are k districts d_1, \dots, d_k . The *social welfare* (or *utility*) that project x_j provides to district d_i is $sw_i(x_j) \in \mathbb{Z}_{\geq 0}$. We assume that utilities are additive; i.e., the utility that an outcome $W \subseteq \mathcal{P}$ provides to district d_i is $sw_i(W) := \sum_{x_j \in W} sw_i(x_j)$. Furthermore, the total social welfare of $W \subseteq \mathcal{P}$ is $sw(W) := \sum_{i \in [k]} sw_i(W)$.

Throughout this work we assume that $sw(x_j)$ and $c(x_j)$ are both $\text{poly}(k, m)$ for each j . (A function f is $\text{poly}(x, y)$ if there exists a $k \geq 0$ such that $f = O((xy)^k)$.) We can relax this assumption using well-known bucketing techniques at the cost of an arbitrarily small ε in the guarantees of our algorithms. See the fully polynomial time approximation scheme for the knapsack problem (Chekuri and Khanna 2005) for an example of this technique.

To model the participatory budgeting setting, we assume that each district deserves some portion of the budget and, in turn, deserves at least the utility it could achieve if it spent its budget on its most preferred projects. Specifically, each district d_i deserves some budget $b_i \geq 0$ where $\sum_i b_i = b$. District d_i deserves utility $f_i := sw_i(W_i)$, where $W_i := \arg \max_{W: c(W) \leq b_i} sw_i(W)$ is d_i 's favorite outcome costing at most b_i .

Definition 1 (District-Fair Outcome). *We say that an outcome W is district-fair (DF) if $sw_i(W) \geq f_i$ for all i .*

Computing f_i is precisely an instance of the knapsack problem; by our assumption that utilities and costs are polynomially bounded, this knapsack instance is solvable in polynomial time (Chekuri and Khanna 2005). Thus, we will assume f_i is known.

Note that the outcome $\bigcup_i W_i$ is both budget-feasible and district-fair, so an outcome with both properties always exists. Our goal is to find a budget-feasible and district-fair outcome

W which maximizes social welfare $sw(W)$. We call our problem *district-fair welfare maximization*. Throughout this paper, we let $W^* := \arg \max_W sw(W)$ be some optimal solution, where the argmax is taken over budget-feasible and district-fair solutions. Similarly, we let $\text{OPT} := sw(W^*)$.

We consider two relaxations of district fairness. The first relaxation extends the concept to *lotteries* over outcomes. We require that each district only needs to be approximately satisfied in expectation. We give an efficient algorithm to compute optimal district-fair lotteries in Section 4.

Definition 2 (ε -District-Fair Lottery). *Given $\varepsilon > 0$, we say that a probability distribution \mathcal{W} over outcomes of cost at most b is an ε -district-fair (ε -DF) lottery if $E_{W \sim \mathcal{W}}[sw_i(W)] \geq f_i - \varepsilon$ for every district d_i .*

The second relaxation is *district-fairness up to one good* (DF1). Intuitively, an allocation is DF1 if each district would be satisfied if one additional project was funded.

Definition 3 (DF1). *An outcome W is DF1 if for every d_i ,*

$$sw_i(W) + \max_{x_j \in (\mathcal{P} \setminus W)} sw_i(x_j) \geq f_i.$$

DF1 is inspired by the well-studied notion of EF1 (envy-freeness up to one good) from the private goods setting (Budish 2011). This relaxation is mild, and unlike relaxations that require district-fairness to hold on average over districts, it is a *uniform* relaxation which provides guarantees for all districts. We study DF1 outcomes in Section 5.

3 NP-Hardness

Our first result shows that the problem of optimizing social welfare subject to district-fairness is NP-hard even in the restricted setting of approval votes (i.e., voters provide binary yes/no opinions over projects) and budgets proportional to district sizes. In fact, our problem remains NP-hard in this restricted setting even when each district contains only one voter and projects have unit costs.

We reduce from exact 3-cover (X3C), which is known to be NP-hard (Garey and Johnson 1979). The idea of our reduction is as follows. Given an instance of X3C, we define a district for each of the elements in the universe, and then add a large amount of dummy districts. We then define a project for each set in our problem instance which gives one utility to the districts corresponding to the elements which it covers. We also define a large set of dummy projects that are approved by all dummy districts. We then ask whether there exists a district-fair outcome that attains high social welfare. An optimal solution for our district-fair welfare maximization problem, then, will first try to solve the X3C instance as efficiently as possible so that it can spend as much of its budget as possible on high-utility dummy projects. We formalize this idea in the following proof.

Theorem 1. *It is NP-complete to decide, given an instance of district-fair welfare maximization and an integer M , whether there exists a budget-feasible and district-fair outcome W such that $sw(W) \geq M$. NP-hardness holds even in the restricted setting of approval votes and budgets proportional to district sizes, and when each district contains one voter and all projects have unit cost.*

Proof. The stated problem is trivially in NP. For NP-hardness we reduce from X3C. In an instance of X3C, we are given a universe $U = \{e_1, \dots, e_{3n}\}$ and a collection $\{S_1, \dots, S_m\}$ of 3-element subsets of U . It is a “yes”-instance if there exists a selection S_{j_1}, \dots, S_{j_n} such that $S_{j_1} \cup \dots \cup S_{j_n} = U$.

Given an instance of X3C, we construct an instance of our problem as follows. Let $M = 3mn + 1$. We have $3n + M$ districts, $D \cup D'$. Let $D = \{d_1, \dots, d_{3n}\}$, where each d_i in D corresponds to element e_i . Additionally, let $D' = \{d_{3n+1}, \dots, d_{3n+M}\}$, where each $d_i \in D'$ is a dummy district. We have $m + 2n + M$ projects, $X \cup X'$. Let $X = \{x_1, \dots, x_m\}$, where $x_j \in X$ corresponds to set S_j , and let $X' = \{x_{m+1}, \dots, x_{m+2n+M}\}$, where each $x_j \in X'$ is a dummy project. Utilities are as follows: every dummy district approves every dummy project, so $\text{sw}_i(x_j) = 1$ for each $i \geq 3n+1$ and $x_j \in X'$. Also, each non-dummy district approves of non-dummy sets to reflect the structure of the X3C instance: that is, for each $i \leq 3n$ we have $\text{sw}_i(x_j) = 1$ if $x_j \in X$ and $e_i \in S_j$. All other utilities are 0: that is, $\text{sw}_i(x_j) = 0$ for all other i and j . Each project has cost 1, and our budget is $b = 3n + M$. We assume all districts contain 1 voter, so $b_i = 1$ for every district d_i . Clearly, $f_i = 1$ for each i . We ask whether there exists a district fair committee with social welfare at least $3n + (2n + M)M$.

If there exists a solution S_{j_1}, \dots, S_{j_n} to the X3C instance, then $W = \{x_{j_1}, \dots, x_{j_n}\} \cup X'$ is an outcome with cost $n + (2n + M) = 3n + M = b$. Clearly, W is district-fair, and its social welfare is $3n + (2n + M)M$, so this is a “yes”-instance for the district-fair welfare-maximization problem.

Conversely suppose that there exists a district-fair budget-feasible outcome W with social welfare at least $3n + (2n + M)M$. Note that all projects in X together give overall welfare at most $3mn < M$. Thus, we must have $X' \subseteq W$ since otherwise the total welfare of W is less than $(2n + M)M$. Hence $|X \cap W| \leq n$. By district-fairness, for each $i = 1, \dots, 3n$, there must be some $x_j \in W$ such that $e_i \in S_j$. These two facts together imply that $\{S_j : x_j \in W\}$ is a solution to the X3C instance. \square

This NP-hardness result holds even if each district consists of a single voter and all projects have unit cost. As we show in Appendix A in the supplementary material, this special case admits a polynomial-time $\frac{1}{2}$ -approximation. Our algorithm is based on a greedy algorithm and a combinatorial argument which “matches away” high utility goods of the optimal solution. One might hope to achieve an approximation result for the general case. A natural approach would be to round the optimal solution to the LP relaxation of the natural ILP formulation of our problem. However, a simple example in Appendix B in the supplementary material shows that the integrality gap of that formulation is unboundedly large, so this approach will not work.

4 Optimal District-Fair Lottery

In this section, we allow randomness and consider lotteries over outcomes. Our main result for the lottery setting is an ε -DF lottery which always achieves the optimal social welfare subject to district fairness. The welfare guarantee is ex post, so that every outcome in the lottery’s support

achieves optimal welfare. For the remainder of this section we let $\varepsilon > 0$ refer to the ε in the ε -DF definition.

Theorem 2. *There is an algorithm which, in poly $(m, k, \frac{1}{\varepsilon})$ time, returns an ε -DF lottery \mathcal{W} such that for all outcomes W in the support of \mathcal{W} , we have $\text{sw}(W) \geq \text{OPT}$.*

The intuition for our algorithm is as follows. We begin by showing that our problem is polynomial-time solvable if the number of districts k is constant. Such an algorithm is useful because we can artificially make the number of districts constant by convexly combining all districts into a single district \bar{d} . We can, then, compute as our solution a utility-optimal outcome W which is fair for \bar{d} but not necessarily fair for each d_i individually. However, we can bias our solution to try and satisfy fairness for certain districts by increasing the weights of these districts in our convex combination. Thus, if W is not fair for d_i , we might naturally increase the proportional share of d_i in the convex combination and recompute W in the hopes that the new outcome we compute will be fair for d_i . We obtain our lottery by repeatedly increasing the weight of districts that do not have their fairness constraint satisfied, and then take a uniform distribution over the resulting outcomes.

Turning to the proof, we begin by describing how to solve our problem in polynomial time when k is a constant. Our algorithm will solve the natural dynamic program (DP). Specifically, consider the true/false value $R(\text{sw}^{(1)}, \dots, \text{sw}^{(k)}, j, b)$ which is the answer to the question, “Does there exist an outcome of cost at most b using projects x_1, x_2, \dots, x_j wherein district d_i achieves social welfare at least $\text{sw}^{(i)}$?” If the answer to this question is yes, then either the desired utilities are possible with the stated budget without using x_j or there is an outcome which uses at most $b - c(x_j)$ budget that doesn’t use x_j in which every district gets at least its specified utility minus how much it values x_j . Thus, $R(\text{sw}^{(1)}, \dots, \text{sw}^{(k)}, j, b)$ is true if and only if either $R(\text{sw}^{(1)}, \dots, \text{sw}^{(k)}, j - 1, b)$ is true or $R(\text{sw}^{(1)} - \text{sw}_1(x_j), \dots, \text{sw}^{(k)} - \text{sw}_k(x_j), j - 1, b - c(x_j))$ is true, giving us a definition by recurrence.

By our assumption that all costs and utilities are polynomially bounded, we can easily solve the dynamic program (DP) for the above recurrence, giving the following result.

Lemma 3. *There is an algorithm that finds a budget-feasible district-fair outcome W with $\text{sw}(W) = \text{OPT}$ in $m^{O(k)}$ time.*

Proof. Our algorithm simply fills in the DP table and returns the outcome corresponding to the entry in our DP table which is true, satisfies $\text{sw}^{(i)} \geq f_i$ for all i and which maximizes $\sum_i \text{sw}^{(i)}$. The recurrence is correct by the above reasoning.

To see why we can fill in the DP table in the stated time, note that we can trivially solve our base case, $R(\text{sw}^{(1)}, \dots, \text{sw}^{(k)}, j, 1)$, for each j and possible value for each $\text{sw}^{(i)}$ in polynomial time. Since $\max_{i,j} \text{sw}_i(x_j)$ is polynomially bounded in m , we need only check polynomially many in m values for each $\text{sw}^{(i)}$. Lastly, since j and b are bounded by a polynomial in m , we conclude that our DP table has $m^{O(k)}$ entries, giving the desired runtime. \square

We now describe our multiplicative-weights-type algorithm to produce our lottery using the above algorithm.⁴ We let $w_i^{(t)} \geq 0$ be the “weight” of district i in iteration t and let $w^{(t)} := \sum_i w_i^{(t)}$ be the total weight in iteration t . Initially, our weights are uniform: $w_i^{(1)} = 1$ for all i .

For any iteration t and district d_i we let $p_i^{(t)} := \frac{w_i^{(t)}}{w^{(t)}}$ be the proportion of the weight that district i has in iteration t . These $p_i^{(t)}$ will induce our convex combination over districts; in particular we let $\tilde{d}^{(t)}$ be a district which values project x_j to extent $\tilde{sw}^{(t)}(x_j) := \sum_i p_i^{(t)} \cdot sw_i(x_j)$ and which deserves $\tilde{f}^{(t)} := \sum_i p_i^{(t)} \cdot f_i$ utility. Also, let sw_{\max} be the maximum welfare of an outcome.

With the above notation in hand, we can give our instantiation of multiplicative weights where $T := \frac{4 \ln k}{\varepsilon^2} \cdot sw_{\max}^2$ is the number of iterations of our algorithm.

1. For all iterations $t \in [T]$:
 - (a) Let W_t be an outcome that maximizes $sw(W_t)$ subject to $\tilde{sw}^{(t)}(W_t) \geq \tilde{f}^{(t)}$ and $c(W_t) \leq b$. We can compute W_t using Lemma 3.
 - (b) Let $m_i^{(t)} := sw_i(W_t) - f_i$ be our “mistakes”, indicating how far off a district was from getting what it deserved.
 - (c) Update weights: $w_i^{(t+1)} \leftarrow w_i^{(t)} \cdot \exp(-\varepsilon m_i^{(t)})$.
2. Return lottery \mathcal{W} , the uniform distribution over $\{W_t\}_{t \leq T}$.

We now restate the usual multiplicative weights guarantee in terms of our algorithm. This lemma guarantees that, on average, the multiplicative weights strategy is competitive with the best “expert.” In the following $\langle p^{(t)}, m^{(t)} \rangle := \sum_i p_i^{(t)} \cdot m_i^{(t)}$ is the usual inner product.

Lemma 4 (Arora, Hazan, and Kale 2012). *For all i we have*

$$\frac{1}{T} \sum_{t \leq T} \langle p^{(t)}, m^{(t)} \rangle \leq \varepsilon + \frac{1}{T} \sum_{t \leq T} m_i^{(t)}.$$

We can use this lemma to show the desired guarantees.

Proof of Theorem 2. We use the algorithm described above.

Our algorithm is polynomial time since it runs for polynomially-many iterations and in each iteration we compute a solution for a problem on only one district which is solvable in polynomial time by Lemma 3. Also, note that by Lemma 3 we know that $c(W_t) \leq b$ for all t , so all outcomes in the lottery are budget-feasible.

We now argue that the above lottery is utility-optimal. Fix an iteration t . Notice that since W^* is fair for all districts then it is fair for $\tilde{d}^{(t)}$. In particular,

$$\tilde{sw}^{(t)}(W^*) = \sum_i p_i \cdot sw_i(W^*) \geq \sum_i p_i f_i = \tilde{f}^{(t)}$$

Thus, W^* is a budget-feasible solution for the problem of finding a max-utility outcome which is fair for $\tilde{d}^{(t)}$. Thus,

⁴We will only need to invoke the above algorithm for the case $k = 1$. This amounts to solving the knapsack problem with a single covering constraint, which to our knowledge is not one of the standard variants of the knapsack problem.

$sw(W_t)$ can only be larger than $sw(W^*)$, meaning that $sw(W_t) \geq \text{OPT}$.

We now argue that the above lottery is ε -DF in expectation. Fix a district d_i . By Lemma 4 we know that

$$\frac{1}{T} \sum_{t \leq T} \langle p^{(t)}, m^{(t)} \rangle \leq \varepsilon + \frac{1}{T} \sum_{t \leq T} m_i^{(t)}. \quad (1)$$

Now notice that by definition of $m_i^{(t)}$ and since our lottery is uniform over all W_t we know that the right-hand-side of Equation (1) is

$$\begin{aligned} \varepsilon + \frac{1}{T} \sum_{t \leq T} m_i^{(t)} &= \varepsilon + \frac{1}{T} \sum_t (sw_i(W_t) - f_i) \\ &= \varepsilon - f_i + \frac{1}{T} \sum_t sw_i(W_t) \\ &= \varepsilon - f_i + \mathbb{E}_{W \sim \mathcal{W}} [sw_i(W)] \end{aligned}$$

Thus, to show that $f_i - \varepsilon \leq \mathbb{E}_{W \sim \mathcal{W}} [sw_i(W)]$, it suffices to show that the left-hand side of Equation (1) is at least 0. That is, we must show $0 \leq \frac{1}{T} \sum_{t \leq T} \langle p^{(t)}, m^{(t)} \rangle$. However, this amounts to simply showing that W_t is fair for $\tilde{d}^{(t)}$; in particular, we have that the left-hand-side is

$$\begin{aligned} \frac{1}{T} \sum_{t \leq T} \langle p^{(t)}, m^{(t)} \rangle &= \frac{1}{T} \sum_{t \leq T} \sum_i p_i^{(t)} \cdot (sw_i(W_t) - f_i) \\ &= \frac{1}{T} \sum_{t \leq T} \tilde{sw}^{(t)}(W_t) - \tilde{f}^{(t)}. \end{aligned}$$

It holds that $\tilde{sw}^{(t)}(W_t) - \tilde{f}^{(t)} \geq 0$ since we always choose a solution which is fair for $\tilde{d}^{(t)}$, and so we conclude that the left-hand-side of Equation (1) is at least 0. \square

5 Optimal DF1 Outcome with Extra Budget

We now study how well we can do if we allow ourselves to overspend the available budget. Certainly it is possible to achieve district fairness and optimal fairness-constrained utility OPT if the algorithm can spend *double* the available budget: we can compute an outcome W_1 with $c(W_1) \leq b$ that is welfare-maximizing without attempting to satisfy district-fairness, and we can compute some outcome W_2 with $c(W_2) \leq b$ that is district-fair (see Section 2); then $W_1 \cup W_2$ satisfies district fairness and we clearly have $c(W_1 \cup W_2) \leq 2b$ and $sw(W_1 \cup W_2) \geq \text{OPT}$. In this section, we show that we can find a solution that requires less than twice the budget, if we slightly relax the district fairness requirement to DF1. Our main result for the DF1 setting shows that, under DF1 fairness, there is a deterministic algorithm which achieves DF1 and optimal social welfare if one overspends a 0.647 fraction of the budget.

Theorem 5. *For any constant $\varepsilon > 0$, there is a poly(m, k)-time algorithm which, given an instance of district-fair welfare maximization, returns an outcome W such that W is DF1, $c(w) \leq (1.647 + \varepsilon)b$, and $sw(W) \geq (1 - \varepsilon)\text{OPT}$.*

Overspending by 64.7% is a worst-case result, and the algorithm may often overspend less. If the context does not permit any overspending, one can run the same algorithm with a reduced budget; then the output will be feasible for the true budget, yet will satisfy weaker fairness and social welfare guarantees. More precisely, given an instance \mathcal{I} and a multiplier $\beta < 1$, we define an instance $\mathcal{I}'(\beta)$, which is identical to \mathcal{I} but in which each district d_i contributes only $\beta \cdot b_i$ and thus deserves utility $f'_i := \text{sw}_i(W'_i)$, where W'_i is d_i 's favorite outcome which costs at most $\beta \cdot b_i$. Additionally, let $\text{OPT}'(\beta)$ represent the maximum achievable social welfare over all district-fair solutions in \mathcal{I}' using a budget of at most $b' := \beta \cdot b$. Then, applying Theorem 5 to $\mathcal{I}'(\beta)$ results in an outcome which is DF1 and utility-optimal on this reduced instance and does not overspend the original budget b .

Corollary 6. *For any constant $\varepsilon > 0$, there is a poly(m, k)-time algorithm which, given an instance \mathcal{I} of district-fair welfare maximization, returns an outcome W such that W is DF1 for $\mathcal{I}'(\frac{1}{1.647})$, $c(W) \leq (1 + \varepsilon)b$, and $\text{sw}(W) \geq (1 - \varepsilon)\text{OPT}'(\frac{1}{1.647})$.*

Our result uses a submodular optimization as a subroutine. If one allows randomization in this subroutine, algorithms with better approximation ratios are known. Thus, we can prove a similar theorem (and corollary) with a *randomized* algorithm which achieves DF1 and optimal social welfare while overspending its budget by only a $\frac{1}{e} \approx .37$ fraction of the budget, with high probability (i.e., with probability $1 - \frac{1}{p(m,k)}$ where $p(m, k)$ is some polynomial in m and k). We defer details of our randomized algorithm to Appendix C in the supplementary material.

In the remainder of this section, we will prove Theorem 5. Our main tool is a notion of the ‘‘coverage’’ of a partial outcome. An outcome has high coverage if we do not need to spend much more money to make it district-fair. On a high level, our proof consists of two main steps. First, we show how to complete an outcome with good coverage into a DF1 outcome. Second, we will show how to frame the problem of finding a solution with good coverage and social welfare as a submodular maximization problem subject to linear constraints, allowing us to use a result by Mizrahi et al. (2018).

We begin by formalizing the coverage of a solution. Roughly, if we imagine that initially every district requires its portion of the budget for fairness, then fractional coverage captures how much less districts must spend to satisfy their own fairness constraints. Thus, if we imagine that our algorithm first spends its budget to satisfy fairness as efficiently as possible, and then spends the remainder of its budget on the highest utility projects, then the coverage of a collection of projects is roughly how much budget this collection ‘‘frees up’’ for the algorithm to spend on the highest utility projects. More formally, we define coverage by way of the notions of fractional outcomes and residual budget requirements.

Definition 4 (fractional outcomes). *A fractional outcome is a vector $p \in \mathbb{R}^m$ where $0 \leq p_j \leq 1$. We overload notation and let the social welfare of p for district d_i be $\text{sw}_i(p) := \sum_j \text{sw}_i(x_j) \cdot p_j$. Similarly the social welfare of p is $\sum_i \text{sw}_i(p)$. Lastly, we define the cost of p as $\sum_j c(x_j) \cdot p_j$.*

We now define the residual budget requirement of a district, given an outcome, which can be understood as the minimum amount of additional money that must be spent to satisfy the district, if fractional outcomes are allowed.

Definition 5 ($\text{resid}_i(W)$). *The residual budget requirement of district d_i given (integral) outcome W is the minimum cost of a fractional outcome p such that $\text{sw}_i(W) + \text{sw}_i(p) \geq f_i$ and $p_j = 0$ for all $x_j \in W$.*

We can now define the coverage of an outcome for a particular district i in terms of the total amount of budget they deserve and their residual budget requirement.

Definition 6 ($\text{cover}_i(W)$). *The coverage of an outcome W for district d_i is the difference between the amount of budget they deserve, b_i , and their residual budget requirement: $\text{cover}_i(W) := b_i - \text{resid}_i(W)$.*

Lastly, we define the coverage of an outcome.

Definition 7 ($\text{cover}(W)$). *The overall coverage of an outcome W is the sum over all districts d_i of the coverage W affords d_i : $\text{cover}(W) := \sum_i \text{cover}_i(W)$.*

Next, we establish a useful property of DF1 solutions. In particular, given a set of projects that achieves relatively good fairness on *average*, we can then buy a small subset of projects that results in fairness up to one good for all districts. In particular, given a collection of projects that covers a $1 - \beta$ fraction of all fairness constraints, we can use at most an extra β fraction of our budget in order to complete this to a DF1 solution. Moreover, this completion is quite intuitive: purchase all projects whose total coverage exceed their cost, until there are no such projects remaining.

Formally, we state the following *DF1 completion* lemma.

Lemma 7 (DF1 Completion). *Given an outcome W with $\text{cover}(W) = b - r$, one can compute in polynomial time a set $W' \supseteq W$ such that W' is DF1 and $c(W') \leq c(W) + r$.*

Proof. We first prove that for every non-DF1 outcome W , there exists a project that we can add to W which increases its coverage by at least $c(x_j)$. Suppose that W is an outcome that fails DF1, and let d_i be a district such that $\text{sw}_i(W) + \text{sw}_i(x_j) < f_i$ for all $x_j \notin W$. Let p be the fractional outcome witnessing $\text{resid}_i(W)$; thus $\text{sw}_i(W) + \text{sw}_i(p) \geq f_i$. We may assume without loss of generality that all but at most one project is integral in p_j (because there is always some optimal p with this property by additivity of sw_i). Since W fails DF1 for d_i , there is some $x_j \notin W$ such that $p(x_j) = 1$. Then $\text{resid}_i(W \cup \{x_j\}) = \text{resid}_i(W) - c(x_j)$ (witnessed by the fractional outcome obtained from p by removing x_j from it). Thus, from definitions, $\text{cover}_i(W \cup \{x_j\}) = \text{cover}_i(W) + c(x_j)$, and hence $\text{cover}(W \cup \{x_j\}) \geq \text{cover}(W) + c(x_j)$.

Now suppose we are given an outcome W with $\text{cover}(W) = b - r$, which fails DF1. We can identify a project x_j as above, add it to W , and increase the coverage by at least $c(x_j)$. We repeat this until the outcome is DF1. This process must stop, since at each step the coverage increases by $c(x_j)$ but by definition the coverage can never exceed b . For the same reason, the cost of the projects we have added to W cannot exceed r , and thus $c(W') \leq c(W) + r$. \square

With this lemma in hand, we now turn to the problem of finding high-coverage outcomes with good welfare. Let $B \geq 0$ be a lower bound on the social welfare we desire. We rephrase our problem as an optimization problem in which we maximize the coverage of an outcome subject to a linear knapsack constraint and a linear covering constraint. The knapsack constraint enforces budget feasibility, and the covering constraint encodes the requirement that the total utility of the outcome is at least B .

$$\begin{aligned} & \max_{W \subseteq \mathcal{P}} \text{cover}(W) \\ & \text{s.t. } \text{sw}(W) \geq B, \\ & \quad c(W) \leq b. \end{aligned} \tag{DF1P}$$

The main tool we apply is a theorem on the maximization of nondecreasing submodular functions of Mizrachi et al. (2018). Recall that a set function is nondecreasing if its value never decreases as elements are added to its input, and submodular if it exhibits diminishing returns.

Definition 8. *Given a finite set Ω , a set function $f : \Omega \rightarrow \mathbb{R}_{\geq 0}$ is nondecreasing and submodular if for every $A, B \subseteq \Omega$ such that $A \subseteq B$ we have $f(A) \leq f(B)$ and $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$ for all $x \in \Omega \setminus B$.*

The theorem we apply is as follows.

Theorem 8 (Mizrachi et al. 2018, Theorem 5). *For each constant $\varepsilon > 0$, there exists a deterministic algorithm for maximizing a nondecreasing submodular function subject to one packing constraint and one covering constraint that runs in time $O(n^{O(1)})$, where $n = |\Omega|$ is the size of the support of the set function, satisfies the covering constraint up to a factor of $1 - \varepsilon$ and the packing constraint up to a factor of $1 + \varepsilon$, and achieves an approximation ratio of 0.353.*

We apply this theorem to find a solution that satisfies a 0.353 fraction of coverage and achieves optimal fairness-constrained utility. Then, we apply Lemma 7 to augment our solution using an additional $1 - 0.353 + \varepsilon$ fraction of our budget in order to obtain a final solution which satisfies full DF1. However, in order to apply Theorem 8, we must first establish that $\text{cover}(W)$ is a nondecreasing submodular function. In particular, note that the coverage functions $\text{cover}_i(W)$ for each district are clearly nondecreasing and submodular. It follows that their sum, $\text{cover}(W)$ is also nondecreasing and submodular, yielding the following lemma.

Lemma 9. *The function $\text{cover}(W)$ is nondecreasing and submodular.*

We are now ready to prove Theorem 5, which applies the DF1 completion lemma to an approximately optimal solution for the problem DF1P.

Proof of Theorem 5. Recall that we have assumed that the maximum utility of an outcome is polynomially bounded in m and k and that the maximum utility is integral. Thus, the value of OPT falls in a polynomial range. For each value B in this range, solve the problem DF1P using the algorithm from Theorem 8. Now consider all values of B for which the algorithm returned a solution with $\text{cover}(W) \geq 0.353b$; such a value must exist since we are guaranteed this condition

when $B = \text{OPT}$ (since for this value, the optimum of problem (DF1P) is b). Among all solutions we found that satisfy $\text{cover}(W) \geq 0.353b$, take the one that maximizes $\text{sw}(W)$. This solution provides social welfare *at least* $(1 - \varepsilon) \text{OPT}$.

We have obtained an outcome W with

$$\text{cover}(W) \geq 0.353b = b - 0.647b,$$

and $\text{sw}(W) \geq (1 - \varepsilon) \text{OPT}$ and $c(W) \leq (1 + \varepsilon)b$. Now apply Lemma 7 to W to obtain a DF1 outcome $W' \supseteq W$ with

$$c(W') \leq c(W) + 0.647b \leq (1 + 0.647 + \varepsilon)b.$$

This outcome W' satisfies the requirements of Theorem 5. \square

6 Discussion

Our results extend to the special case of unit costs, also known as committee selection. In committee selection, we elect a committee to represent voters in a larger governmental body such as a parliament. Often, to ensure local representation, the electorate is split into voting districts, which elect their representatives separately. The districts may be apportioned different numbers of representatives, for example based on district size. While this scheme guarantees each district representation, it may well be possible to increase the welfare of the voters in a district, for example by electing a diverse array of candidates with expertise in various areas who can gather votes from across the electorate. Thus, it is natural for all districts to elect the committee together if we impose district-fairness constraints. This way, we can maximize social welfare of the final committee while guaranteeing each district fair representation. This gives a more holistic view of committee selection in exactly the same way we addressed participatory budgeting, only instead of pooling the budget between districts, we now pool seats on a committee.

Our model implicitly treats districts as atoms, and so district fairness is a kind of *individual rationality* property. In turn, individual rationality is a type of strategyproofness: it incentivizes districts not to leave the central election and instead hold a separate one. Is it possible to design a voting scheme that is fully strategyproof for districts, so that districts do not have incentives to misreport the utilities of their residents? Unfortunately not: Peters (2018) proves an impossibility theorem about committee elections which implies that there does not exist a voting rule that is efficient, district-fair, and also strategyproof. This result holds even for approval votes.

Several open questions remain. Most obvious is the question of whether can we achieve welfare maximization and DF1 in polynomial time while guaranteeing to overspend the budget by less than $1/e$. More broadly, it would be interesting to study our problem with more general utility functions such as submodular or even general monotone valuation functions. Additionally, it would be exciting to study approximation algorithms which promise full district fairness. In Appendix B in the supplementary material, we present an algorithm which satisfies district fairness and provides a $1/2$ -approximation to optimal district-fair social welfare in the special case of unanimous districts; it would be interesting to extend this result to the general case.

References

- Arora, S.; Hazan, E.; and Kale, S. 2012. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing* 8(1): 121–164.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48(2): 461–485.
- Aziz, H.; Elkind, E.; Huang, S.; Lackner, M.; Sánchez-Fernández, L.; and Skowron, P. 2018. On the Complexity of Extended and Proportional Justified Representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 902–909.
- Aziz, H.; and Shah, N. 2020. Participatory Budgeting: Models and Approaches. In Rudas, T.; and Péli, G., eds., *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer.
- Brill, M.; Freeman, R.; Janson, S.; and Lackner, M. 2017. Phragmén’s Voting Methods and Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 406–413.
- Brown, J. R. 1979. The knapsack sharing problem. *Operations Research* 27(2): 341–355.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6): 1061–1103.
- Cabannes, Y. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization* 16(1): 27–46.
- Cabannes, Y. 2014. Contribution of Participatory Budgeting to provision and management of basic services. *London: IIED*.
- Cabannes, Y. 2017. Participatory budgeting in Paris: Act, reflect, grow. *Another city is possible with participatory budgeting* 179–203.
- Chekuri, C.; and Khanna, S. 2005. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM Journal on Computing* 35(3): 713–728.
- Cheng, Y.; Jiang, Z.; Munagala, K.; and Wang, K. 2019. Group fairness in committee selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM EC)*, 263–279.
- 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 (WINE)*, 384–399.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM EC)*, 575–592. Extended version arXiv:1805.03164.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 2.
- Fujimoto, M.; and Yamada, T. 2006. An exact algorithm for the knapsack sharing problem with common items. *European Journal of Operational Research* 171(2): 693–707.
- Garey, M. R.; and Johnson, D. S. 1979. *Computers and intractability*, volume 174. Freeman San Francisco.
- Goel, A.; Krishnaswamy, A. K.; Sakshuwong, S.; and Aitamurto, T. 2019. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)* 7(2): 1–27.
- Hifi, M.; M’Halla, H.; and Sadfi, S. 2005. An exact algorithm for the knapsack sharing problem. *Computers & Operations Research* 32(5): 1311–1324.
- Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301): 13–30.
- Janson, S. 2016. Phragmén’s and Thiele’s election methods. Technical report. ArXiv:1611.08826 [math.HO].
- Jiang, Z.; Munagala, K.; and Wang, K. 2020. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 463–472.
- Mizrachi, E.; Schwartz, R.; Spoerhase, J.; and Uniyal, S. 2018. A tight approximation for submodular maximization with mixed packing and covering constraints. arXiv:1804.10947.
- Peters, D. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 1549–1557.
- Peters, D.; and Skowron, P. 2020. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM EC)*, 793–794.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J. A.; Basanta Val, P.; and Skowron, P. 2017. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 670–676.
- Talmon, N.; and Faliszewski, P. 2019. A framework for approval-based budgeting methods. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, volume 33, 2181–2188.
- Williamson, D. P.; and Shmoys, D. B. 2011. *The design of approximation algorithms*. Cambridge University Press.
- Yamada, T.; and Futakawa, M. 1997. Heuristic and reduction algorithms for the knapsack sharing problem. *Computers & operations research* 24(10): 961–967.
- Yamada, T.; Futakawa, M.; and Kataoka, S. 1998. Some exact algorithms for the knapsack sharing problem. *European Journal of Operational Research* 106(1): 177–183.

Appendix: District-Fair Participatory Budgeting

A $\frac{1}{2}$ -Approximation for Unanimous Districts with Unit Costs

In this section we study approximation algorithms for the simplest version of our problem which we know to be NP-hard: when each districts consist of a single voter and every project has unit cost. In fact, we will study a strictly more general setting than each district consisting of a single voter; namely, we study the setting where each district is “unanimous.” Formally, we study instances of district-fair welfare maximization where $c(x_j) = 1$ for all j and $\text{sw}_i(x_j) \in \{0, |d_i|\}$ for all i where $|d_i|$ is the number of voters in d_i . For this setting we will give a $\frac{1}{2}$ -approximation.

Our algorithm will make use of the following notion of conditional coverage which builds on Definition 6.

Definition 9. *The coverage of a project x_j given an outcome W is $\text{cover}(x_j|W) := \sum_i \text{cover}_i(W \cup \{x_j\}) - \text{cover}_i(W)$.*

Notice that by our assumption of unit cost and unanimous districts we have that $\text{cover}(x_j|W) \in \mathbb{Z}_{\geq 0}$.

We now present our greedy algorithm that satisfies district fairness and achieves a $1/2$ -approximation to the optimal district-fair utility for the setting of unanimous districts and unit costs. Formally, the algorithm, which we call the *Unanimous Greedy Algorithm (UGA)* proceeds as follows.

1. Given an instance \mathcal{I} , initialize $W_0 \leftarrow \emptyset$.
2. For $j \in [b]$:
 - (a) Let $c_j := \max_{x_j} \text{cover}(x_j|W_{j-1})$ be the max possible coverage and let $X_j := \{x_j : \text{cover}(x_j|W_{j-1}) = c_j\}$ be all projects which achieve this coverage.
 - (b) Let $x_j := \arg \max_{x_j \in X_j} \text{sw}(x_j)$ be the max covering project with maximum utility.
 - (c) Update $W_j \leftarrow W_{j-1} \cup \{x_j\}$.
3. Return W_b .

Theorem 10. *Given an instance \mathcal{I} consisting of unanimous approval districts, UGA returns a solution which satisfies district fairness and achieves a $1/2$ -approximation to the optimal district-fair utility.*

Proof. Let W represent the result of UGA, and let W^* represent the optimal district-fair outcome. Furthermore, let N_2, N_1 and N_0 be all projects purchased by UGA that had conditional coverage at least 2, exactly 1 and exactly 0 when purchased by UGA respectively. Clearly W is district-fair and budget-feasible and so we need only argue that it achieves at least $\text{sw}(W^*)/2$ utility.

Now, consider the following subproblem, which we will call \mathcal{I}' , which intuitively is our original instance \mathcal{I} but where all of N_2 is forced to be in a solution and no projects from N_0 are available. More formally, \mathcal{I}' is \mathcal{I} but where our budget is changed to $b' := b - |N_2| - |N_0|$, f_i is changed to $f'_i := \max(f_i - \text{sw}_i(N_2), 0)$ for all i and the set of purchasable projects is $\mathcal{P}' := \mathcal{P} \setminus N_2$. $\text{sw}_i(x_j)$ is the same for all i, j in

\mathcal{I}' as in \mathcal{I} . Notice that the coverage of any project in \mathcal{I}' is at most 1 but the total coverage required for fairness is b' , meaning that every budget-feasible and district-fair solution for \mathcal{I}' has size exactly b' . Also notice that N_1 is not only feasible for \mathcal{I}' but also attains the optimal utility among all district-fair and budget-feasible solutions.

We claim that there exists a subset $N_1^* \subseteq W^* \setminus N_2$ which is district-fair and budget-feasible for \mathcal{I}' . To see this, note that we can iteratively build N_1^* by initializing it to \emptyset and then repeatedly adding to it any $x_j \in W^* \setminus N_1^* \setminus N_2$ such that $\text{cover}(x_j|N_1^*)$ in \mathcal{I}' is at least 1. After b' such additions we are guaranteed to have a district-fair and budget-feasible solution for \mathcal{I}' and such an x_j always exists since $W^* \setminus N_2$ is district-fair for \mathcal{I}' . As noted above, any district-fair and budget-feasible solution for \mathcal{I}' has size b' and so $|N_1^*| = b'$.

Thus, since N_1 is optimal for \mathcal{I}' , we know that N_1 achieves at least as high utility as N_1^* , i.e., $|N_1| = |N_1^*| = b'$, and

$$\text{sw}(N_1^*) \leq \text{sw}(N_1). \quad (2)$$

It remains to understand the utility of $W^* \setminus N_1^*$. However, note that at least half of the projects other than N_1 must be in the N_0 phase. That is, $|N_0| \geq |N_2|/2$. Intuitively, this means that UGA “frees up” at least $\frac{b-b'}{2}$ money to spend on high-utility projects. Let $\mathcal{P}_0 := \mathcal{P} \setminus \{N_2 \cup N_1\}$ be all projects not in N_2 or N_1 . We have that $\text{sw}(N_0)$ is the utility of the top $\frac{b-b'}{2}$ projects in \mathcal{P}_0 . On the other hand, consider projects in $W^* \setminus N_1^*$. We can divide these into projects which are in N_2 and N_1 and those which are not. In particular, let $W_{1,2}^* := W^* \setminus N_1^* \cap (N_1 \cup N_2)$ and let $W_0^* := W^* \setminus N_1^* \cap \mathcal{P}_0$ so that $W_{1,2}^* \cup W_0^* = W^* \setminus N_1^*$. Now notice that trivially

$$\text{sw}(W_{1,2}^*) \leq \text{sw}(N_1) + \text{sw}(N_2). \quad (3)$$

On the other hand, $|W_0^*| = b - b' - |W_{1,2}^*| \leq b - b'$ and $W_0^* \subseteq \mathcal{P}_0$ and so $\text{sw}(W_0^*)$ is at most the utility of the $b - b'$ highest utility projects in \mathcal{P}_0 . Since our utilities are additive and $\text{sw}(N_0)$ is the utility of the $\frac{b-b'}{2}$ highest utility projects in \mathcal{P}_0 , it follows that

$$\text{sw}(W_0^*) \leq 2 \cdot \text{sw}(N_0). \quad (4)$$

Since $W^* = N_1^* \cup W_0^* \cup W_{1,2}^*$, we can combine the above bounds to conclude our $\frac{1}{2}$ -approximation. Namely, applying the additivity of our utilities and combining Equations 2, 3 and 4 we have

$$\begin{aligned} \text{sw}(W^*) &= \text{sw}(N_1^*) + \text{sw}(W_0^*) + \text{sw}(W_{1,2}^*) \\ &\leq 2 \cdot \text{sw}(N_0) + 2 \cdot \text{sw}(N_1) + \text{sw}(N_2) \\ &\leq 2 \cdot \text{sw}(W) \end{aligned}$$

and so we conclude that $\text{sw}(W) \geq \frac{\text{sw}(W^*)}{2}$. \square

B Integrality Gap

Here, we investigate the integrality gap of the natural LP for our problem. As a reminder, the integrality gap of an LP measures how much better a fractional solution can do than an integral solution. An unbounded integrality gap shows that any analysis of an approximation algorithm which charges

the value of its integral solution to the value of the optimal LP gives an unboundedly-bad approximation ratio. For this reason integrality gaps are sometimes taken as evidence of hardness of approximation. For more details on the topic of integrality gaps see Williamson and Shmoys (2011).

We will show that our LP has an unbounded integrality gap which suggests that approximation algorithms which return budget-feasible and district-fair solutions with nearly-optimal social welfare may be difficult or impossible to attain for the general case.

Formally our LP and its integrality gap are as follows. Our LP has a variable y_j for each project x_j corresponding to the extent to which we choose x_j .

$$\begin{aligned} & \max \sum_j y_j \cdot \text{sw}(x_j) \\ & \text{s.t.} \sum_j y_j \cdot c(x_j) \leq b \\ & \sum_j y_j \cdot \text{sw}_i(x_j) \geq f_i \quad \forall i \\ & 0 \leq y_j \leq 1 \quad \forall j \end{aligned} \tag{DFLP}$$

We let $DFLP(\mathcal{I})$ correspond to the polytope corresponding to the above LP for an instance \mathcal{I} of district-fair welfare maximization.

The integrality gap of DFLP is defined as

$$\min_{\mathcal{I}} \frac{\max_{y \in DFLP(\mathcal{I}) \cap \mathbb{Z}^m} \sum_j y_j \cdot \text{sw}(x_j)}{\max_{y \in DFLP(\mathcal{I})} \sum_j y_j \cdot \text{sw}(x_j)}.$$

The basic idea of our integrality gap construction is as follows. We will construct an instance of social-welfare maximization where the preferences of each district are ‘‘circular’’. In particular, each district will like two projects and every project will be liked by exactly two districts. As in our NP-hardness proof, we will also have a collection of dummy projects which are given very high utility by dummy districts which deserve no utility. An optimal fractional solution will be able to choose each non-dummy project to extent essentially $\frac{1}{2}$ to satisfy district-fairness and then spend its remaining budget on high-utility dummy projects. On the other hand, the optimal integral solution will have to spend its entire budget satisfying fairness.

Theorem 11. *There does not exist a function f such that the integrality gap of DFLP is at most $\frac{1}{f(k,m)}$. Further, this integrality gap holds even when all projects have unit cost.*

Proof. Fix $k \in \mathbb{Z}_{\geq 1}$ and a sufficiently small $\varepsilon > 0$. We define our instance of social-welfare maximization on k districts where d_1, d_2, \dots, d_{k-1} will be non-dummy districts and the district d_k will be a dummy district. Similarly, we will have $2(k-1)$ projects where x_1, x_2, \dots, x_{k-1} will be non-dummy projects and the remaining projects $x_k, \dots, x_{2(k-1)}$ will be dummy projects.

For each non-dummy district d_i we let $b_i = 1$ and define its utility for project x_j as

$$\text{sw}_i(x_j) := \begin{cases} 1 + \varepsilon & \text{if } j = i \\ 1 & \text{if } j = (i + 1 \bmod k - 1) + 1 \\ 0 & \text{otherwise} \end{cases}$$

For the dummy district d_k we let $b_k = 0$ and define its utility for project x_j as

$$\text{sw}_k(x_j) := \begin{cases} B & \text{if } x_j \text{ is a dummy project} \\ 0 & \text{otherwise} \end{cases}$$

for B sufficiently large to be chosen later. Notice that $\text{sw}(x_j) = B$ for each dummy project x_j . Lastly, we let our budget $b = k - 1$ and we let $c(x_j) = 1$ for all x_j .

Now notice that each non-dummy district d_i has $f_i = 1 + \varepsilon$. Consequently, any district-fair integral solution must include all non-dummy projects, namely x_1, x_2, \dots, x_{k-1} . However, since $b = k - 1$, it follows that the only district-fair integral solution is $W_{\text{int}} := \{x_1, x_2, \dots, x_{k-1}\}$ where $\text{sw}(W_{\text{int}}) = (2 + \varepsilon)(k - 1)$.

On the other hand, consider the following fractional solution y . For each non-dummy project x_j we let $y_j = \frac{1+\varepsilon}{2}$. For each dummy project x_j we let y_j be $(1 - \frac{1+\varepsilon}{2})$. Clearly $\sum_j y_j \leq b$. Moreover, notice that for each district d_i we have $\sum_j y_j \cdot \text{sw}_i(x_j) = \frac{1+\varepsilon}{2}(2 + \varepsilon) \geq 1 + \varepsilon = f_i$ and so our solution is indeed in the polytope of DFLP. However, since $y_j = (1 - \frac{1+\varepsilon}{2})$ for each dummy project we have that $\sum_j y_j \cdot \text{sw}(x_j) \geq B(k - 1)(1 - \frac{1+\varepsilon}{2})$.

Thus, for the above instance we have that the ratio of the optimal integral solution to the optimal fractional solution is at most

$$\frac{(2 + \varepsilon)(k - 1)}{B(k - 1)(1 - \frac{1+\varepsilon}{2})} \leq \frac{10}{B}.$$

Since B can be chosen independently of k and m , we have that the above instance has integrality gap strictly less than $\frac{1}{f(k,m)}$ for any function f of k and m . \square

We note that the proof of the above result also rules out any integrality gap which is which is larger than $o(\frac{1}{c})$ where c is the total number of voters across all districts.

C Randomized Optimal DF1 Outcome with Extra Budget

In this section we give our randomized analogues of Theorem 5 and Corollary 6. We use the notation of Section 5 throughout this section. Whereas our deterministic algorithms overspend budget by 64.7%, our randomized algorithms will only overspend it by $\frac{1}{e}$ with high probability. Formally, we show the following theorem.

Theorem 12. *There is a poly(m, k)-time algorithm which, given an instance of district-fair welfare maximization, returns an outcome W such that W is DF1, $c(W) \leq (1 + \frac{1}{e} + \varepsilon)b \approx 1.37b$ with high probability, and $\text{sw}(W) \geq (1 - \varepsilon)\text{OPT}$ for any fixed constant $\varepsilon > 0$.*

As with Corollary 6 for Theorem 5, we immediately have a corollary which gives an algorithm which does not overspend its budget.

Corollary 13. *There is a poly(m, k)-time algorithm which, given an instance of district-fair welfare maximization, returns an outcome W such that W is DF1 for $\mathcal{I}' \left(\frac{1}{1+1/e} \right)$,*

$c(W) \leq (1 + \varepsilon)b$ with high probability, and $\text{sw}(W) \geq (1 - \varepsilon)\text{OPT}' \left(\frac{1}{1 + 1/e} \right)$ for any fixed constant $\varepsilon > 0$.

On a high level, this proof will closely follow that of Theorem 5. In particular, it uses the same submodular optimization framing of the problem (i.e., DF1P). However, there are some notable differences. In particular, we leverage the following randomized result from Mizrachi et al. (2018) instead of the previous deterministic result from Mizrachi et al. (2018).

Theorem 14 (Mizrachi et al. 2018, Theorem 1). *For each constant $\varepsilon > 0$, there exists a randomized algorithm for maximizing a nondecreasing submodular function subject to one packing constraint and one covering constraint that runs in time $O(|\Omega|^{O(1)})$, satisfies the packing constraint, satisfies the covering constraint up to a factor of $1 - \varepsilon$, and achieves an expected approximation ratio of $1 - \frac{1}{e} - \varepsilon$.*

Additionally, we require Hoeffding's inequality (Hoeffding 1963), which bounds the probability that the sum of a sequence of independent random variables deviates from its expectation; we restate it below.

Theorem 15 (Hoeffding's Inequality). *Given a sequence of n independent random variables Y_1, \dots, Y_n , where each Y_i takes a value in the range $[\alpha_i, \beta_i]$, we have that*

$$\Pr \left(\mathbb{E} \left[\sum_{k=1}^n Y_k \right] - \sum_{k=1}^n Y_k \geq t \right) \leq e^{-\frac{2t^2}{\sum_{i=1}^n (\beta_i - \alpha_i)^2}},$$

where $t \geq 0$.

We now present the main lemma of the section, which states that there exists a polynomial-time algorithm which, given a guess of the optimal district-fair social welfare $B \leq \text{OPT}$, returns an outcome W that is DF1 fair, overspends the budget by approximately b/e with high probability, and achieves at least $(1 - \varepsilon)B$ social welfare.

Lemma 16. *Given an instance of DF1P, there is an algorithm that runs in polynomial time which returns an outcome W such that W is DF1, $c(W) \leq (1 + \frac{1}{e} + \varepsilon)b$ with high probability, and $\text{sw}(W) \geq (1 - \varepsilon)B$, where $\varepsilon > 0$ is an arbitrary constant and $B \leq \text{OPT}$.*

Proof. Note that, by Lemma 7, it suffices to find an outcome W' such that $c(W') \leq b$, $\text{cover}(W') \geq (1 - \frac{1}{e} - \varepsilon)b$, and $\text{sw}(W') \geq (1 - \varepsilon)B$, as we can complete this solution into a DF1 solution with $(\frac{1}{e} + \varepsilon)b$ more budget.

Let $\varepsilon_0 := \varepsilon/2$. By Lemma 9, we may apply Theorem 14 in order to find a solution W such that $\text{cover}(W) \geq (1 - \varepsilon_0)(1 - 1/e)$ in expectation, $\text{sw}(W) \geq B$, and $c(W) \leq b$. We now show how to transform this guarantee in expectation into one with high probability by using Hoeffding's inequality and an averaging argument.

In order to apply Hoeffding's inequality to our setting, let Y_j represent the coverage of the j^{th} run of our application of Theorem 14. We know that each run is independent, and therefore the Y_j 's are also independent. Let $n = \omega(\log k / \varepsilon_0^2)$, and define a sequence of n independent random variables Y_1, \dots, Y_n representing the coverage of n runs of the mechanism. Furthermore, because coverage is bounded between 0 and b , we have that $\beta_i = b$ and $\alpha_i = 0$ for all $i \in [n]$.

By Hoeffding, we have

$$\begin{aligned} \Pr \left(\mathbb{E} \left[\sum_{j=1}^n Y_j \right] - \sum_{j=1}^n Y_j \geq \varepsilon_0 \left(1 - \frac{1}{e} \right) bn \right) \\ \leq \exp \left(-\frac{2(\varepsilon_0 (1 - \frac{1}{e}) bn)^2}{\sum_{i=1}^n (\beta_i - \alpha_i)^2} \right) \\ = \exp \left(-2\varepsilon_0^2 \left(1 - \frac{1}{e} \right)^2 n \right). \end{aligned}$$

Because we set $n = \omega(\log k / \varepsilon_0^2)$, this probability goes to 0 polynomially quickly. Therefore, we know that, with high probability, $\mathbb{E} \left[\sum_{j=1}^n Y_j \right] - \sum_{j=1}^n Y_j \leq \varepsilon_0 (1 - \frac{1}{e}) bn$, or $\sum_{j=1}^n Y_j \geq bn(1 - \varepsilon_0) (1 - \frac{1}{e}) - \varepsilon_0 (1 - \frac{1}{e}) bn$. By an averaging argument, this means that, with high probability, there exists a Y_i such that $Y_i \geq (1 - \varepsilon_0 - \varepsilon_0) (1 - \frac{1}{e}) b = (1 - \varepsilon) (1 - \frac{1}{e}) b$, as desired. \square

With this lemma in hand, we are ready to prove Theorem 12.

Proof of Theorem 12. Recall that we have assumed that the maximum utility of an outcome is polynomially bounded in m and k and that the maximum utility is integral. Thus, the value of OPT falls in a polynomial range. For each value B in this range, run the procedure from Lemma 16. By a union bound over these polynomially-many applications of Lemma 16, we have that all applications of Lemma 16 with $B \leq \text{OPT}$ succeed with high probability, resulting in a collection of solutions; one for each B . Now consider all values of B for which the algorithm returned a solution with $c(W) \leq (1 + \frac{1}{e} + \varepsilon)b$; such a value must exist since with high probability we are guaranteed this condition when $B = \text{OPT}$ and we know that, with high probability, all applications of Lemma 16 succeed for all $B \leq \text{OPT}$. Among all solutions we found that satisfy $c(W) \leq (1 + \frac{1}{e} + \varepsilon)b$, take the one that maximizes $\text{sw}(W)$. This solution provides utility at least OPT . \square