

Asymptotically Fair and Truthful Allocation of Public Goods

POUYA KANANIAN, University of Waterloo, Canada

SEYED MAJID ZAHEDI, University of Waterloo, Canada

We study the fair and truthful allocation of m divisible public items among n agents, each with distinct preferences for the items. To aggregate agents' preferences fairly, we follow the literature on the fair allocation of public goods and aim to find a core solution. For divisible items, a core solution always exists and can be calculated efficiently by maximizing the Nash welfare objective. However, such a solution is easily manipulated; agents might have incentives to misreport their preferences. To mitigate this, the current state-of-the-art finds an approximate core solution with high probability while ensuring approximate truthfulness. However, this approach has two main limitations. First, due to several approximations, the approximation error in the core could grow with n , resulting in a non-asymptotic core solution. This limitation is particularly significant as public-good allocation mechanisms are frequently applied in scenarios involving a large number of agents, such as the allocation of public tax funds for municipal projects. Second, implementing the current approach for practical applications proves to be a highly nontrivial task. To address these limitations, we introduce PPGA, a (differentially) Private Public-Good Allocation algorithm, and show that it attains asymptotic truthfulness and finds an asymptotic core solution with high probability. Additionally, to demonstrate the practical applicability of our algorithm, we implement PPGA and empirically study its properties using municipal participatory budgeting data.

1 INTRODUCTION

Unlike the allocation of private goods, where each item goes to a single agent, public goods allow multiple agents to benefit from an allocated item. In this paper, we study the problem of fairly allocating m divisible public goods among n agents in a truthful manner. Different agents hold distinct preferences for the items. Each item has a size, and the total size of allocated items should not exceed the available capacity. The fair allocation of divisible public goods is a fundamental problem in social choice theory with many real-world applications. Examples include: (1) budget allocations between services such as healthcare, education, and defense at a state or national level; (2) municipal budget allocations to improve utilities such as libraries, parks, gyms, and roads; (3) shared memory allocations between files with different sizes; and (4) time allocations between activities during an event.

An allocation mechanism produces outcomes based on reported preferences of all agents. Agents need not reveal their true preferences but strategically report them to maximize their utility. For instance, consider a setting where there are one or more commonly preferred items. Such items are highly likely to be allocated regardless of the reported preferences of a single agent. Given this and assuming that other agents report their preference truthfully, agents could be incentivized to *free-ride* by falsely claiming disinterest in commonly preferred items and reporting preferences only for their individually preferred items. By doing so, free riders increase the chances of their individually preferred items being allocated under a fair allocation mechanism.

To aggregate agents' preferences fairly, we focus on the classic game theoretic notion of the core [5, 16]. The core generalizes well-studied notions of proportionality and Pareto efficiency by ensuring group-wise fairness, providing fair outcomes to each agent subset relative to its size. The notion of the core has been extensively studied in the context of public-good allocation [9, 10, 13, 31, 32]. For allocating divisible public goods, the core always exists, and it can be efficiently calculated by maximizing *Nash welfare* (NW) objective (i.e., the product of agents' utilities) [9].

However, the core outcome is easy to manipulate; agents might have incentives to misreport their preferences.

To address this issue, Fain et al. [9] propose a method that aims to find an approximate core solution with high probability while also achieving approximate truthfulness. This method relies on the exponential mechanism derived from differential privacy [29]. The exponential mechanism uses a scoring function to assign a score to each outcome. Subsequently, a sample is drawn from a distribution that exponentially weights outcomes based on their scores. This guarantees that the selected outcome’s score is approximately maximized with high probability.

Differentially private mechanisms are approximately dominant-strategy truthful, as the general guarantee of differential privacy ensures that no agent can significantly alter the outcome of the mechanism by unilaterally changing their reported input. In the context of the exponential mechanism, the level of differential privacy—and consequently, truthfulness—is contingent on the sensitivity¹ of the scoring function to the reported input of any individual agent. Higher sensitivity corresponds to a lower quality of the guarantee.

The use of the exponential mechanism for public-good allocation faces two primary challenges. Firstly, while the NW objective appears to be an ideal choice for the scoring function, its direct use is hindered by its high sensitivity to each agent’s reported preferences. This limitation arises because the NW objective is not separable². To overcome this, Fain et al. [9] propose a novel proxy function as a substitute for the NW objective in the scoring function.

The introduced proxy function strikes a balance between reducing the sensitivity of the scoring function to enhance the approximation to truthfulness and retaining sufficient sensitivity to ensure an acceptable approximation to the core. However, the adoption of the proxy function introduces an approximation error in the core that could grow with an increasing number of agents, resulting in a non-asymptotic core solution. This limitation is particularly significant as public-good allocation mechanisms are frequently applied in scenarios involving a large number of agents, such as participatory budgeting elections for distributing municipal budgets.

Secondly, sampling an m -dimensional allocation from a distribution poses a significant practical challenge. To tackle this, Fain et al. [9] propose employing the hit-and-run method [39] to sample an allocation from an “approximately right distribution.” However, implementing the hit-and-run method for practical applications proves to be a highly nontrivial task (as discussed in the conclusion of Sec. 2.2 in [28]). Moreover, the implications of the extra approximation on the guarantees of truthfulness and core remain unclear.

1.1 Our Contributions

In Sec. 4, we introduce PPGA, a novel differentially private algorithm for public-good allocation. A key feature of PPGA is its approach to maximize the NW objective in a differentially private way without requiring a proxy objective. As previously discussed, the non-separable nature of the NW objective poses challenges in deploying differentially private mechanisms [9]. To tackle this challenge, we employ a key technique called *global variable consensus optimization* [1]. Consensus transforms the NW objective into a separable form that splits easily. Leveraging the *alternating direction method of multipliers (ADMM)* [15, 17] enables us to maximize the NW objective in a distributed manner. And this further allows us to employ the *Gaussian mechanism* [29] from differential privacy to achieve truthfulness.

¹Informally, the sensitivity of a function is the maximum change in its output resulting from a change in its input.

²A function f is separable with respect to a partition of a variable x into n sub-vectors $x = (x_1, \dots, x_n)$ if $f(x) = \sum f_i(x_i)$.

In Sec. 5, we analytically study the properties of our proposed algorithm. Our primary technical contribution lies in demonstrating that PPGA attains asymptotic truthfulness and finds an asymptotic core solution with high probability. To our knowledge, PPGA is the first polynomial-time algorithm that provides such guarantees.

In Sec. 6, we demonstrate that PPGA can be deployed in practice to solve large-scale public-good allocation problems. To this end, we implement PPGA and utilize our implementation to compare the outcome of PPGA with a core solution using data obtained from real-world participatory budgeting elections conducted in various cities in Poland [40]. The code for our implementation is provided at <https://github.com/uwaterloo-mast/PPGA>.

2 RELATED WORKS

Fair resource allocation. There is an extensive literature on designing mechanisms for allocation of private goods without money, which is often referred to as cake cutting [35]. For public goods, the fair allocation problem has been studied in the context of fair public decision-making [3], multi-agent knapsack problems [12], multi-winner elections [32], and participatory budgeting [33]. Truthful aggregation of agents' preferences has also been studied for public decision-making [14, 18, 36]. However, the settings in these works are unrelated to ours as they focus on maximizing social welfare (i.e., the sum of agents' utilities), whereas our focus is on maximizing Nash welfare.

The work most closely related to this paper is that of Fain et al. [9]³, which finds an approximate core solution with high probability while achieving approximate truthfulness. However, due to its reliance on several approximations, their approach fails to produce an asymptotic core solution. As the number of agents increases, their method's approximation error for fairness (core) could grow. In contrast, our approximation guarantee does not suffer from this issue. By combining the Gaussian mechanism with ADMM to directly optimize the NW objective, our method provides asymptotic truthfulness and finds an asymptotic core solution with high probability.

Differentially private convex optimization. In recent years, differentially private convex programming has been utilized to allocate private goods [4, 21, 22, 25, 26]. These methods often employ the dual ascent technique as a key tool [1]. The dual ascent method involves a sequence of two updates: the primal update optimizes the Lagrangian while fixing the dual variable, and the dual update takes a gradient ascent step to update the dual variable given the optimized primal variable. The dual ascent method cannot be used for maximizing the NW objective. This is because, as we show in Sec. 4, the Lagrangian for the convex program is an affine function of some components of the primal variable. This makes the primal update fail as the dual problem is unbounded below for most values of the dual variable [1]. We avoid this by optimizing the augmented Lagrangian instead of the Lagrangian.

Differentially private ADMM. It is well known that in ADMM, the objective function's value converges to the optimal solution at a rate of $O(1/K)$, where K denotes the number of iterations [19]. Recent years have witnessed significant research efforts aimed at merging differential privacy and ADMM [23, 24, 27, 38, 42, 43]. Although related, our work differentiates itself from these works in several aspects. Firstly, while previous studies focus on the convergence rate of the objective function, we study the convergence of a primal variable to an approximate core solution. To the best of our knowledge, our work is first to prove an asymptotic, game-theoretic property for a primal variable within differentially private ADMM. Secondly, unlike prior work that introduces noise to the local variables, PPGA adds noise to the global variable (as detailed in Sec. 4). Finally,

³Their notion of core is based on capacity, wherein a blocking coalition receives a proportional share of the capacity instead of a proportional share of utility (refer to Def. 2).

many studies on differentially private ADMM rely on a restrictive assumption regarding the strong convexity of the objective function. This assumption, however, does not hold for the NW objective.

3 PRELIMINARIES

In this section, we first formally define the public-good allocation problem and its desired properties. We then overview differential privacy as a tool for mechanism design.

A detailed summary of our notations is presented in Appx. A

3.1 Problem Formulation

We consider a public-good allocation problem with n agents and m divisible public items ($m \ll n$). The size of each item j is denoted by $s_j \in \mathbb{R}_{>0}$, and the size vector is denoted by $s = (s_1, \dots, s_m)$. The total available capacity is $c \in \mathbb{R}_{>0}$. An allocation is a vector $z = (z_1, \dots, z_m) \in [0, 1]^m$, where z_j represents the fraction of the total capacity that is allocated to item j . The set of all feasible allocations is denoted by:

$$Z = \{z \in [0, 1]^m \mid \|z\|_1 \leq 1, cz \leq s\}.$$

Agent i 's utility function for allocation $z \in Z$ is denoted by $U_i(z)$ and is parameterized by the utility vector $u_i = (u_{i1}, \dots, u_{im})$. In this paper, we consider a subclass of utility functions that are differentiable, strictly increasing, concave, and L -Lipschitz continuous:

$$|U_i(z) - U_i(z')| \leq L\|z - z'\|_2.$$

A notable example of such utility function is the *linear* utility function: $U_i(z) = \sum_{j=1}^m u_{ij}z_j$. Sigmoid and hyperbolic tangent are two other examples. Without loss of generality, we assume that $U_i(z) \in [0, 1]$ for all i and $z \in Z$, and that $u_i \in U = [0, 1]^m$ for every i .

3.2 Mechanism Design for Public Goods

A randomized allocation mechanism M produces a probability distribution over feasible allocations given agents' reported utilities $u = (u_1, \dots, u_n) \in U^n$. Agents need not report their true utilities. They report their utilities strategically to optimize their total utility possibly taking into account what (they think) other agents report. If agents are always incentivized to report their true utilities, no matter what others do, then the mechanism is *truthful*:

DEFINITION 1 ((DOMINANT-STRATEGY) TRUTHFULNESS). Let U_i be agent i 's utility function parameterized by i 's true utility vector u_i . A randomized mechanism M is (ϵ, δ) -truthful if for every i , $u'_i \in U$, and $u_{-i} \in U^{n-1}$, we have⁴:

$$\mathbb{E}_{z \sim M(u_i, u_{-i})} [U_i(z)] \geq (1 - \epsilon) \mathbb{E}_{z \sim M(u'_i, u_{-i})} [U_i(z)] - \delta.$$

If $\epsilon, \delta = 0$, then M is *exactly truthful*. Approximate truthfulness is desirable in settings in which the approximation parameters ϵ and δ tend to 0 as the number of agents n grows large. We call this *asymptotic truthful*. We next define the classic notion of the *core*:

DEFINITION 2 (CORE). A set of agents A form a *blocking coalition* if there exists an allocation $z' \in Z$ such that $(|A|/n)U_i(z') \geq U_i(z)$ for every agent $i \in A$ with at least one strict inequality. An allocation is a *core outcome* if it admits no blocking coalitions.

In this definition, when a subset A of agents deviates, they can choose any feasible allocation with the full capacity c . However, their utility is scaled down by a factor of $|A|/n$. An alternative way of defining a core solution is where a deviating coalition A could choose any allocation with a capacity of $c|A|/n$ instead of c , but their utilities would not be scaled down [13, 37]. For $|A| = n$,

⁴Subscript $-i$ is used to refer to all agents other than agent i .

both notions capture Pareto efficiency. However, for $|A| = 1$, each notion provides a different definition of proportionality: one based on utility and one based on capacity.

For divisible goods, a core solution is guaranteed to exist, and it coincides with the *max Nash welfare* (MNW) solution⁵:

LEMMA 3 (MNW $\hat{=}$ CORE). *Suppose that U_i is differentiable, strictly increasing, and concave for all agents i . The allocation that maximizes $\sum_i \log(U_i(z))$ subject to $z \in Z$ is a core solution of the public-good allocation problem⁶.*

This lemma shows that the exact MNW solution is a core solution. However, if the optimization problem is solved approximately, then the solution is not guaranteed to be a core solution. For such cases, we need an approximation to the core that can still provide a meaningful guarantee.

DEFINITION 4 (APPROXIMATE CORE). *For $\epsilon, \delta \geq 0$, an allocation $z \in Z$ is an (ϵ, δ) -core outcome if there exists no set of agents $A \subseteq N$ and no allocation $z' \in Z$ such that:*

$$(|A|/n)U_i(z') \geq (1 + \epsilon)U_i(z) + \delta$$

for all $i \in A$ with at least one strict inequality.

When ϵ and δ converge asymptotically to 0 as n grows large, we call this *asymptotic core*. We now provide a lemma to show that approximate MNW implies approximate core.

LEMMA 5 (APPROXIMATE MNW $\hat{=}$ CORE). *Let $\epsilon, \delta \geq 0$ and $z \in Z$ be an allocation that satisfies:*

$$\frac{1}{n} \sum_i \frac{U_i(z')}{U_i(z) + \delta/(1 + \epsilon)} \leq 1 + \epsilon \quad (1)$$

for any $z' \in Z$. Then z is an (ϵ, δ) -core outcome.

PROOF. Suppose for contradiction that z is not an (ϵ, δ) -core outcome. Then there exist a set A and an allocation z' s.t. $(|A|/n)U_i(z') \geq (1 + \epsilon)U_i(z) + \delta \forall i \in A$ with at least one strict inequality. This implies: $(1/n) \sum_{i \in A} U_i(z)/(U_i(z) + \delta/(1 + \epsilon)) > 1 + \epsilon$, which contradicts (1). \square

3.3 Mechanism Design via Differential Privacy

Our goal is to design a mechanism that attains both approximate truthfulness and core. This can be accomplished by pursuing an approximate MNW solution in a differentially private (DP) manner, as DP inherently implies approximate truthfulness (see Lemma 13). Informally, a mechanism satisfies DP if its output is almost equally likely to be observed on every *adjacent* inputs. Inputs are deemed adjacent if they vary in only one element. In the domain of mechanisms for public-good allocation, inputs represent reported utilities from all agents. Therefore, $u, u' \in U^n$ are regarded as adjacent if they differ solely in the reported utility of a single agent. We now proceed to present the formal definition of DP.

DEFINITION 6 (DP [7]). *A randomized mechanism M is (ϵ, δ) -DP if for any pair of adjacent inputs $u, u' \in U^n$ and every subset of outputs $O \in Z$, it satisfies:*

$$\mathbb{P}[M(u) \in O] \leq e^\epsilon \mathbb{P}[M(u') \in O] + \delta.$$

⁵Fain et al. [9] present a version of this lemma for homogeneous functions of degree 1 (Corollary 2.3). For completeness, we provide a proof for the lemma in Appx. C.1.

⁶In this paper, all logarithms are natural.

Since the adjacency relation is symmetric, we further have:

$$\begin{aligned}\mathbb{P}[M(u) \in O] &\geq e^{-\epsilon} \mathbb{P}[M(u') \in O] - e^{-\epsilon} \delta \\ &\geq e^{-\epsilon} \mathbb{P}[M(u') \in O] - \delta.\end{aligned}$$

Here, ϵ and δ control the desired level of privacy. In general, smaller values provide stronger privacy guarantees but result in higher levels of noise being required to be injected, which can adversely affect the quality of the output. When $\delta = 0$, M satisfies the standard ϵ -DP.

DP is preserved by *post-processing*, meaning the DP guarantee of a mechanism does not diminish by manipulating its output. In essence, if M is (ϵ, δ) -DP, then applying a randomized mapping f to $M(u)$ retains the (ϵ, δ) -DP property as well ([8, Proposition 2.1]). *Rényi differential privacy (RDP)* is a relaxation of DP:

DEFINITION 7 (RDP [30]). *A randomized mechanism M is (α, ϵ) -RDP with order $\alpha > 1$ if for any two adjacent inputs $u, u' \in U^n$, it satisfies: $D_\alpha(M(u) \| M(u')) \leq \epsilon$, where D_α is the Rényi divergence of order α defined as:*

$$D_\alpha(P \| Q) \triangleq \frac{1}{\alpha - 1} \log \left(\mathbb{E}_{x \sim Q} \left[\left(\frac{P(x)}{Q(x)} \right)^\alpha \right] \right).$$

RDP allows for tighter analysis of composition, a key property that enables the modular construction of DP algorithms. If M_1 satisfies ϵ_1 -DP and M_2 satisfies ϵ_2 -DP, then simultaneously releasing the outputs of M_1 and M_2 guarantees $(\epsilon_1 + \epsilon_2)$ -DP. Similar guarantee holds for RDP:

LEMMA 8 (RDP ADDITIVITY [30]). *Let M_1 and M_2 be (α, ϵ_1) -RDP and (α, ϵ_2) -RDP, respectively. Then $M_{1,2}$ defined as $M_{1,2}(u) \triangleq (M_1(u), M_2(u))$ is $(\alpha, \epsilon_1 + \epsilon_2)$ -RDP.*

This guarantee holds even when M_2 is chosen *adaptively* based on the output of M_1 . Lemma 8 expresses the intuitive concept of a *privacy budget*. The cumulative privacy loss during the execution of an iterative mechanism can be easily tracked using the RDP's additivity property.

The *Gaussian mechanism* is a tool to achieve RDP. To provide a formal definition, we first define the ℓ_2 sensitivity.

DEFINITION 9 (L2 SENSITIVITY). *Let $f : U^n \mapsto \mathbb{R}^m$ be an m -dimensional function. The ℓ_2 sensitivity of f is defined as:*

$$\Delta_2(f) \triangleq \max_{u, u' \in U^n} \|f(u) - f(u')\|_2,$$

where the max is taken over all adjacent inputs u and u' .

The Gaussian mechanism computes a vector-valued function and perturbs each coordinate with noise drawn from a Gaussian distribution. The magnitude of the noise is adjusted based on the ℓ_2 sensitivity of the function.

DEFINITION 10 (GAUSSIAN MECHANISM). *Let $f : U^n \mapsto \mathbb{R}^m$ be an m -dimensional function with an ℓ_2 sensitivity of $\Delta_2(f)$. Denoting a multivariate normal distribution with mean vector μ and covariance matrix Σ as $\mathcal{N}(\mu, \Sigma)$, the Gaussian mechanism $M_{f, \alpha, \epsilon}^G$ is defined as follows for $\alpha > 1$ and $\epsilon > 0$:*

$$M_{f, \alpha, \epsilon}^G(u) \triangleq \mathcal{N}(f(u), \sigma^2 I_m),$$

where I_m is the $m \times m$ identity matrix, and $\sigma^2 = \alpha \Delta_2^2(f) / 2\epsilon$.

LEMMA 11 (GAUSSIAN MECHANISM $\hat{=}$ RDP [30]). *Let $f : U^n \mapsto \mathbb{R}^m$ be a vector-valued function with an ℓ_2 sensitivity of $\Delta_2(f)$. Then $M_{f, \alpha, \epsilon}^G$ is (α, ϵ) -RDP.*

Finally, RDP implies DP:

LEMMA 12 (RDP TO DP [30]). *If M is (α, ϵ) -RDP, M is $(\epsilon + \log(1/\delta)/(\alpha - 1), \delta)$ -DP for any $\delta \in (0, 1)$.*

And DP implies approximate truthfulness.

LEMMA 13 (DP TO TRUTHFULNESS). *Let M be (ϵ, δ) -DP for some $\epsilon, \delta < 1$. Then M is (ϵ, δ) -truthful.*

PROOF. Consider any agent i , and let $U_i : Z \mapsto [0, 1]$ be agent i 's utility parameterized according to their true utility vector u_i . Define the set $S(t) = \{z \mid U_i(z) > t\}$. Since M is (ϵ, δ) -DP, for any $u = (u_i, u_{-i}) \in U^n$ and $u'_i \in U$, the following inequality holds:

$$\mathbb{P}[M(u) \in S(t)] \geq e^{-\epsilon} \mathbb{P}[M(u'_i, u_{-i}) \in S(t)] - \delta. \quad (2)$$

Given the definition of $S(t)$, we can rewrite (2) as:

$$\mathbb{P}[U_i(M(u)) > t] \geq e^{-\epsilon} \mathbb{P}[U_i(M(u'_i, u_{-i})) > t] - \delta. \quad (3)$$

Now, considering any random variable $X \in [0, 1]$, we know that $\mathbb{E}[X] = \int_0^1 \mathbb{P}[X > t] dt$. Integrating over both sides of (3) yields:

$$\begin{aligned} \mathbb{E}_{z \sim M(u)}[U_i(z)] &\geq e^{-\epsilon} \mathbb{E}_{z \sim M(u'_i, u_{-i})}[U_i(z)] - \delta \\ &\geq (1 - \epsilon) \mathbb{E}_{z \sim M(u'_i, u_{-i})}[U_i(z)] - \delta, \end{aligned}$$

where the second inequality follows because $e^{-\epsilon} \geq 1 - \epsilon$. \square

4 ALGORITHM

In this section, we present the main artifact of our work, PPGA, an algorithm that directly maximizes the NW objective in a DP manner. Our approach involves a transformation of the NW objective into a separable form. Initially, we reframe the optimization problem of Lemma 3 into a consensus problem. Next, we convert the consensus problem into a distributed optimization using ADMM. Finally, to insure truthfulness, we deploy the Gaussian mechanism.

4.1 Distributed Maximization of Nash Welfare

Consider the NW objective function $\theta(z) = \sum_i \theta_i(z) = \sum_i \log(U_i(z))$. This function does not split as z is shared across terms. To make the objective *separable*, we rewrite the program with local variables x_i and a global variable z :

$$\begin{aligned} \text{Max.} \quad & \sum_{i=1} \theta_i(x_i), \\ \text{s.t.} \quad & z = x_i \quad \forall i \in 1, \dots, n, \\ & x_i \in Z \quad \forall i \in 1, \dots, n. \end{aligned} \quad (4)$$

This is referred to as the *global variable consensus problem*, as it requires all local variables to reach agreement by being equal. Consensus transforms the additive objective, which does not split, into a separable objective, which splits easily. Given the new separable objective, we next apply ADMM to solve the optimization problem in a distributed way. To this end, we first construct the partial augmented Lagrangian [20, 34] for (4):

$$\begin{aligned} L^\rho(x, z, \gamma) &= \sum_i L_i^\rho(x_i, z, \gamma_i) \\ &= \sum_i \left(\theta_i(x_i) - \gamma_i^T (x_i - z) - \frac{\rho}{2} \|x_i - z\|_2^2 \right), \end{aligned}$$

where γ_i is a dual variable corresponding to the constraint $z = x_i$, and $\rho > 0$ is a penalty parameter. Note that L^ρ is also separable in $x = (x_1, \dots, x_n)$ and splits into n separate functions L_i^ρ . Given this,

ADMM can be used to *efficiently* solve the convex program. ADMM is an iterative algorithm which consists of the following iterations:

$$x_i^{(k)} := \operatorname{argmax}_{x \in Z} (L_i^\rho(x, z^{(k-1)}, y_i^{(k-1)})) \quad \forall i \in 1, \dots, n, \quad (5a)$$

$$z^{(k)} := \operatorname{argmax}_z (L^\rho(x^{(k)}, z, y^{(k-1)})), \quad (5b)$$

$$y_i^{(k)} := y_i^{(k-1)} + \rho(x_i^{(k)} - z^{(k)}) \quad \forall i \in 1, \dots, n. \quad (5c)$$

Since L_i^ρ 's are separate functions in x , we can solve (5a) for each $x_i^{(k)}$ separately in parallel. We can also derive closed-form solution to (5b) by setting $\partial L^\rho / \partial z = \sum_i (y_i^{(k-1)} + \rho(x_i^{(k)} - z^{(k)})) = 0$, which implies:

$$z^{(k)} = \frac{1}{n} \sum_i x_i^{(k)} + \frac{1}{n\rho} \sum_i y_i^{(k-1)}. \quad (6)$$

ADMM is guaranteed to find the MNW solution [1]. This means that ADMM is guaranteed to produce a core solution. However, ADMM is not truthful. To address this issue, we next modify ADMM to design a mechanism that guarantees DP (and consequently truthful).

4.2 DP for Maximizing Nash Welfare

To illustrate our proposed mechanism, it might be beneficial to interpret ADMM as an interactive process. At iteration k , each agent i calculates local variable $x_i^{(k)}$ autonomously. Given $z^{(k-1)}$ and $y_i^{(k-1)}$, the value of $x_i^{(k)}$ depends solely on agent i 's own utility. With $x_i^{(k)}$ and $z^{(k)}$ known, each agent i independently calculates $y_i^{(k)}$. These local variables are then submitted by agents, aggregated by the algorithm, and used to compute the global variable $z^{(k)}$. This resultant global variable is broadcast back to the agents for the next iteration.

Now, in ensuring DP, it is imperative that the global variable's value remains insensitive to any individual local variable. To achieve this, we employ the Gaussian mechanism, adding a normal random vector $q^{(k)}$ to $z^{(k)}$:

$$z^{(k)} = \frac{1}{n} \sum_i x_i^{(k)} + \frac{1}{n\rho} \sum_i y_i^{(k-1)} + q^{(k)}. \quad (7)$$

According to (5c), we have:

$$\sum_i y_i^{(k)} = \sum_i (y_i^{(k-1)} + \rho(x_i^{(k)} - z^{(k)})). \quad (8)$$

Replacing $z^{(k)}$ from (7) into (8), we get $\sum_i y_i^{(k)} = -\rho n q^{(k)}$. Using this, we can rewrite (7) as:

$$z^{(k)} = \frac{1}{n} \sum_i x_i^{(k)} - q^{(k-1)} + q^{(k)}. \quad (9)$$

This update rule shows how $z^{(k)}$ can be calculated by adding Gaussian noise to the average of $x_i^{(k)}$'s. The magnitude of the noise can be adjusted to achieve a desired DP guarantee.

Alg. 1 shows the pseudocode of our proposed (differentially) private public-good allocation mechanism, PPGA. The algorithm takes as parameters K , ϵ , δ , and α . The value of K dictates the number of iterations that the algorithm performs. The remaining parameters— ϵ , δ , and α —set the desired level of privacy guarantee. Specifically, ϵ and δ directly control the level of DP (and consequently the truthfulness) that the algorithm ensures (see Thm. 15). Lastly, α governs the variance of the noise according to the Gaussian mechanism.

Algorithm 1: Private public-good allocation (PPGA)

```

1 Parameters:  $K \in \mathbb{Z}$ ,  $\epsilon, \delta \in (0, 1)$ ,  $\alpha > 1$ 
2  $\epsilon' \leftarrow (1/K)(\epsilon - \log(1/\delta)/(\alpha - 1))$ ;
3  $\sigma^2 \leftarrow \alpha/n^2\epsilon'$ ;
4  $q^{(0)}, z^{(0)}, y_i^{(0)}, x_i^{(0)} = \mathbf{0}_m \forall i \in 1, \dots, n$ ;
5 for  $k = 1, \dots, K$  do
6    $x_i^{(k)} \leftarrow \operatorname{argmax}_{x_i \in Z} (L_i^\rho(x_i, z^{(k-1)}, y_i^{(k-1)})) \forall i \in 1, \dots, n$ ;
7    $q^{(k)} \sim \mathcal{N}(0, \sigma^2 I_m)$ ;
8    $z^{(k)} \leftarrow (1/n) \sum_i x_i^{(k)} + q^{(k)} - q^{(k-1)}$ ;
9    $y_i^{(k)} \leftarrow y_i^{(k-1)} + \rho(x_i^{(k)} - z^{(k)}) \forall i \in 1, \dots, n$ ;
10 end
11  $\bar{z} \leftarrow (1/K) \sum_{k=1}^K z^{(k)}$ ;
12  $\hat{z} \leftarrow \Pi_Z(\bar{z})$ ;
13 Output:  $\hat{z}$ 

```

At each iteration k , each agent i privately calculates their optimal allocation $x_i^{(k)}$ given $y_i^{(k-1)}$ and $z^{(k-1)}$. This step can be efficiently executed in parallel for all agents. The algorithm then publicly releases $z^{(k)}$ as a noisy average of $x_i^{(k)}$'s. Given $z^{(k)}$ and $x_i^{(k)}$, each agent i then privately calculates $y_i^{(k)}$ for the next iteration. After K iterations, the algorithm calculates \bar{z} , the time average of $z^{(k)}$'s and returns \hat{z} , the Euclidean projection of \bar{z} onto Z^7 .

Remarks. The integration of DP into ADMM inherently presents a trade-off between accuracy and privacy (truthfulness). Achieving a more accurate MNW solution requires a higher number of iterations. Fixing the amount of privacy loss per iteration, a higher number of iterations means a higher cumulative privacy loss, resulting in a weaker privacy guarantee. On the other hand, achieving a stronger privacy guarantee requires a lower cumulative privacy loss. Fixing the number of iterations, a lower cumulative privacy loss means a higher level of noise per iteration, resulting in diminished accuracy. In PPGA, the expected value of the noise magnitude at each iteration is:

$$\mathbb{E} \left[\|q^{(k)}\|_2^2 \right] = m\sigma^2 = \frac{K m \alpha}{n^2(\epsilon - \log(1/\delta)/(\alpha - 1))}.$$

As a final remark, we note that our proposed algorithms is not actually online or interactive. All computations are performed by the algorithm itself rather than by the agents. The agents submit their private utility vectors and at the end observe a final allocation vector. As we prove in Sec. 5, our algorithm guarantees DP. As a result, agents are assured that their private information is not revealed to other agents. Our mechanism also guarantees asymptotic truthfulness, which means that as n increases, agents do not have any incentives to misreport their private utilities.

5 ANALYSIS

In this section, we first show that PPGA guarantees asymptotic truthfulness. We then show that our mechanism produces an asymptotic core solution with high probability. As discussed in Sec. 4, to establish asymptotic properties of PPGA, we set $K = \Theta(n)$, $\epsilon = \Theta(1/\log(n))$, $\delta = \Theta(1/\sqrt{n})$, $\alpha = \Theta(\log^2(n))$, and $m = o(\sqrt{n})$. All omitted proofs are provided in Appx. C.

⁷ $\Pi_Z(z) = \operatorname{argmin}_{z' \in Z} \|z - z'\|_2^2$.

5.1 Asymptotic truthfulness

To analyze the end-to-end privacy guarantee of Alg. 1, we separately analyze the DP guarantee of each iteration. Leveraging the properties of the Gaussian mechanism, we show that each iteration of the algorithm ensures (α, ϵ') -RDP. With the additivity property of RDP (Lemma 8), after K iterations, Alg. 1 achieves $(\alpha, K\epsilon')$ -RDP. It then easily follows from Lemma 12 that PPGA is (ϵ, δ) -DP.

LEMMA 14. *Alg. 1 is (ϵ, δ) -DP.*

PROOF. Alg. 1 is a composition of K iterations. At each iteration k , the private data is $x^{(k)}$, and the publicly released data is $z^{(k)}$. Note that $\gamma^{(k)}$ is not publicly released as each $\gamma_i^{(k)}$ is privately calculated for each agent i . The z -update step at Line 8 of Alg. 1 is a direct application of the Gaussian mechanism with vector-valued function $f(x) = \frac{1}{n} \sum_i x_i$. Let x and x' be two adjacent inputs that are identical except in their i th element, $x_i \neq x'_i$. Then $\|f(x) - f(x')\|_2 = \frac{1}{n} \|x_i - x'_i\|_2$. Since $x_i, x'_i \in [0, 1]^m$ and $\|x_i\|_1, \|x'_i\|_1 \leq 1$, we have $\|x_i - x'_i\|_2 \leq (\|x_i\|_2^2 + \|x'_i\|_2^2)^{1/2} \leq \sqrt{2}$. This means $\Delta_2(f) \leq \sqrt{2}/n$. Therefore, it follows from Lemma 11 that each iteration k of the algorithm is (α, ϵ') -RDP. Consequently, according to Lemma 8, the composition of the K iterations satisfies $(\alpha, \bar{\epsilon})$ -RDP, where $\bar{\epsilon} = K\epsilon' = \epsilon - \log(1/\delta)/(\alpha - 1)$. The calculation of \bar{z} and projecting it into Z are merely post-processing steps. Since DP is immune to post-processing, it follows from Lemma 12 that Alg. 1 is (ϵ, δ) -DP \square

We next establish our first technical result:

THEOREM 15. *Alg. 1 is asymptotically truthful.*

PROOF. Since Alg. 1 is (ϵ, δ) -DP, it follows directly from Lemma 13 that it is (ϵ, δ) -truthful. Given that $\delta = \Theta(1/\sqrt{n})$ and $\epsilon = \Theta(1/\log(n))$, Alg. 1 is asymptotically truthful. \square

5.2 Asymptotic Core

Let $x = (x_1, \dots, x_n)$, $\gamma = (\gamma_1, \dots, \gamma_n)$, $w = (x, z, \gamma) \in W \triangleq (Z^n, \mathbb{R}^m, \mathbb{R}^{mn})$, and define $w^{(k)} = (x^{(k)}, z^{(k)}, \gamma^{(k)})$. To show that \hat{z} is an approximate core solution, it is essential to derive an upper bound on $\max_{z \in Z} \sum_i \frac{U_i(z)}{U_i(\bar{z})}$. To achieve this, we initially derive an upper bound on $\max_{z \in Z} \sum_i \frac{U_i(z)}{U_i(\bar{x})}$, where $\bar{x} = \frac{1}{K} \sum_k x^{(k)}$. Subsequently, we establish an upper bound on the distance between \hat{z} and any \bar{x}_i . Leveraging the L -Lipschitz continuity of $U_i(x)$, we utilize Lemma 5 to establish \hat{z} as an approximate core solution:

LEMMA 16. *Let $\{w^{(k)}\}$ and $\{q^{(k)}\}$ be sequences generated by Alg. 1. Then for any $z \in Z$, we have:*

$$\frac{1}{n} \sum_i \frac{U_i(z)}{U_i(\bar{x}_i)} \leq 1 + \frac{\rho}{K} \sum_{k=1}^K (z^{(k)} - z)^T q^{(k)} + \frac{\rho}{2K}. \quad (10)$$

Lemma 16 provides an upper bound on $\sum_i \frac{U_i(z)}{U_i(\bar{x}_i)}$. Yet, to show that Alg. 1 finds an approximate core solution, we need an upper bound on $\sum_i \frac{U_i(z)}{U_i(\hat{z})}$. We achieve this by upper-bounding the distance between \hat{z} and any \bar{x}_i :

LEMMA 17. *Let $\{w^{(k)}\}$ and $\{q^{(k)}\}$ be sequences generated by Alg. 1. Let $w^* = (x^*, z^*, \gamma^*)$ be the solution to (4), with $x_i^* = z^*$ for all i . Then we have:*

$$\|\bar{x} - G\hat{z}\|_2 \leq \frac{1}{\rho K} \|\gamma^*\|_2 + \frac{1}{\rho K} \left(2n\rho^2 \sum_{k=1}^K (z^{(k)} - z^*)^T q^{(k)} + n\rho^2 \|z^*\|_2^2 + \|\gamma^*\|_2^2 \right)^{\frac{1}{2}}. \quad (11)$$

Finally, we establish PPGA’s asymptotic fairness.

THEOREM 18. *The outcome of Alg. 1 is an asymptotic core solution with probability $1 - \frac{1}{n} - \frac{1}{n^m}$.*

6 EXPERIMENTS

In this section, we aim to show that PPGA can be deployed in practice to solve large-scale public-good allocation problems. To this end, we implement Alg. 1 in Python using CVXPY, an open-source Python-embedded modeling language for convex optimization problems [6]. PPGA is highly parallelizable, particularly in the concurrent computation of x and y for all agents. We leverage this feature in our implementation by parallelizing the execution of the code that computes x and y across agents, distributing the computational workload across multiple processes using Python’s multiprocessing package. The code for our implementation is provided at <https://github.com/uwaterloo-mast/PPGA>.

To conduct experiments, we leverage real-world data from Pabulib.org, an open participatory budgeting library [40]. Our experiments focus on 12 election instances, selected primarily based on the size of their voter population and the average number of votes per voter⁸. Each instance involves a collection of projects with associated costs and a designated total budget. Voters express their preferences for the projects by casting approval votes for one or more projects. We summarize the key characteristics of these election instances in Appx. B, and full details of each instance, such as project costs, are provided with our code (located in the `final_data` folder).

As just mentioned, all election instances involve approval votes and indivisible projects. We utilized these instances to obtain new instances in which agents have cardinal utilities, and fractional allocations are acceptable. To transform approval votes into cardinal utilities, we adopt the (randomized) *cost-utility* approach [11] using the following procedure: For each voter i and project j , if i does not approve j , we set $u_{ij} = 0$. However, if i does approve j , we sample u_{ij} uniformly from the interval $[0.85, 1.15]$. This method ensures that the utility of each voter is approximately proportional to the budget allocated to the projects they support⁹.

In the concluding remarks of Sec. 4.2, we discuss the impact of parameter values on the added noise. In the opening paragraph of Sec. 5, we establish guidelines for these parameters to guarantee our asymptotic properties. There are also established practical norms for acceptable ϵ and δ values in differential privacy. We follow these norms and instantiate the parameters accordingly. Specifically, we set $\epsilon = c_\epsilon / \log(n)$, $\delta = c_\delta / \sqrt{n}$, and $K = c_K n$, where $c_\epsilon = 1.5$, $c_\delta = 0.3$, and $c_K = 0.001$. We further set α such that $\log(1/\delta)/(\alpha - 1) = \epsilon/2$. This way, values for ϵ and δ approximate 0.3 and 0.001, respectively, keeping the noise magnitude ($\mathbb{E} [\|q^{(k)}\|_2^2]$) under $3e-4$ for the majority of instances.

To empirically study properties of PPGA, we utilize our implementation and compare the outcome of PPGA with a core solution for each instance. To find a core solution, we solve the convex optimization of Lemma 3 by running Alg. 1 without adding noise. We compare our results using the following metrics:

- **Social welfare (SW).** We calculate the social welfare for an allocation z as $(1/n) \sum_i U_i(z)$. SW serves as an indicator of the overall satisfaction achieved collectively by all agents from the allocation.
- **Proportionality score (PS).** We define the proportionality score of voter i for an allocation z as $U_i(z) / \max_{z' \in Z} U_i(z')$. PS evaluates whether each voter receives their fair share relative to what

⁸At the time of our selection, there were about 60 instances with more than 10k votes, many of which with only a single vote per voter.

⁹Suppose that the set of projects supported by agent i is denoted as P_i . Then, i ’s utility is modeled as $U_i(z) = \sum_{j \in P_i} u_{ij} z_j$, where $u_{ij} \sim U(0.85, 1.15)$ for $j \in P_i$. Here, for any $z \in Z$, z_j (where $z_j \leq s_j/c$) represents the fraction of the total budget allocated to project j . Therefore, i ’s utility is roughly proportional to the budget allocated to projects supported by i .

| Inst. | Core's PS | | PPGA's PS | | SD ($\div m$) |
|-------|--------------------|------|--------------------|------|--------------------|
| | Min ($\times n$) | Avg | Min ($\times n$) | Avg | |
| 1 | 91.06 | 0.27 | 88.10 | 0.27 | 0.00007 |
| 2 | 236.8 | 0.29 | 9.411 | 0.28 | 0.00016 |
| 3 | 235.1 | 0.18 | 157.7 | 0.17 | 0.00014 |
| 4 | 216.5 | 0.38 | 62.47 | 0.38 | 0.00022 |
| 5 | 15.01 | 0.32 | 15.10 | 0.32 | 0.00009 |
| 6 | 246.0 | 0.38 | 36.47 | 0.38 | 0.00030 |
| 7 | 11.05 | 0.28 | 10.49 | 0.28 | 0.00044 |
| 8 | 122.6 | 0.32 | 127.4 | 0.32 | 0.00007 |
| 9 | 163.4 | 0.33 | 163.2 | 0.33 | 0.00002 |
| 10 | 152.5 | 0.16 | 116.9 | 0.15 | 0.00033 |
| 11 | 519.8 | 0.44 | 503.4 | 0.44 | 0.00002 |
| 12 | 261.2 | 0.57 | 69.07 | 0.56 | 0.00003 |

Table 1. Proportionality score and statistical distance.




Fig. 1. Social welfare of PPGA normalize to that of core.

they could achieve in the best-case scenario. This is captured in the notion of proportionality, which is satisfied if there is no other allocation that improves the utility of at least one voter by a factor greater than n ($|A| = 1$ in Def. 2). In other words, to satisfy proportionality, the PS value should be $\geq 1/n$ for all voters, or equivalently, the minimum value of PS across voters multiplied by n should be ≥ 1 . We report both the minimum (multiplied by n) and the average of PS values across all voters.

- **Statistical distance (SD).** We measure the distance between z and a core solution z^* by computing their *total variation distance*, defined as $(1/2)\|z - z^*\|_1$. This distance, which also serves as a *metric*¹⁰, quantifies the statistical proximity between the allocation z and a core solution z^* . Two allocations over m items are considered statistically close if their total variation distance is a negligible function in m . To facilitate comparison, we report the normalized total variation distance by dividing it by m .

Since PPGA is a randomized algorithm, we report the average value of each metric over 50 runs of PPGA for each instance.

Fig. 1 illustrates the social welfare under PPGA normalized to that under the core solution, while Tab. 1 summarizes proportionality scores and statistical distances across all election instances.

¹⁰A metric on a set satisfies: (1) non-negativity, (2) identity of indiscernibles, (3) symmetry, and (4) the triangle inequality.

These results uncover several crucial insights. Firstly, the statistical distance between the budget allocation under PPGA and the core solution remains consistently close to zero in all instances, hovering below 0.0004 for all cases. Secondly, the observed discrepancy in social welfare values between PPGA and the core solution consistently falls below 3% across all election instances. Lastly, the minimum PS value $\times n$ exceeds 1 for all instance, indicating that PPGA satisfies the proportionality criteria for all instances. The average PS values tend to be slightly higher under the core solution for some instances, but the discrepancy between the average PS values under PPGA and the core remains below 2% in all instances. Collectively, these findings strongly signify the high level of fairness achieved by PPGA.

Remarks. We note that the primary objective of our experiments is not to empirically validate the properties of PPGA. Rather, our goal is to demonstrate the practical applicability of PPGA to large-scale public-good allocation problems. However, we anticipate our results to be robust for any other linear utility models. This stems from our proof in Thm. 18, where we demonstrate that the distance between \bar{x} and \hat{z} asymptotically approaches zero with high probability. For any linear utility model, it can be shown that the distance between \bar{x} and z^* also asymptotically approaches zero with high probability. Consequently, we would expect the statistical distance between \hat{z} and z^* to be negligible for any linear utility model. For other concave utility models, the statistical distance between \hat{z} and z^* might be slightly higher, depending on the curvature of the function. Nevertheless, one can demonstrate that the difference in the value of the Nash-welfare objective for \hat{z} and z^* asymptotically approaches zero with high probability, implying similar results for PS.

7 CONCLUSION

In this paper, we introduce PPGA, a mechanism designed for the fair and truthful allocation of divisible public goods. PPGA achieves fairness by directly maximizing the NW objective and ensures truthfulness by deploying the Gaussian mechanism from differential privacy. We showed that PPGA is asymptotically truthful and finds an asymptotic core solution with high probability. By conducting experiments using real-world data from participatory budgeting elections, we showcased the practical applicability of PPGA.

REFERENCES

- [1] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning* 3, 1 (2011), 1–122.
- [2] Stephen P Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge University Press.
- [3] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 629–646.
- [4] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2015. Privacy and truthful equilibrium selection for aggregative games. In *Proceedings of the International Conference on Web and Internet Economics (WINE)*. 286–299.
- [5] Gerard Debreu and Herbert Scarf. 1963. A limit theorem on the core of an economy. *International Economic Review* 4, 3 (1963), 235–246.
- [6] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research* 17, 1 (2016), 2909–2913.
- [7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography (TCC)*. 265–284.
- [8] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Theoretical Computer Science* 9, 3-4 (2014), 211–407.
- [9] Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The core of the participatory budgeting problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*. 384–399.
- [10] Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 575–592.

- [11] Piotr Faliszewski, Jarosław Flis, Dominik Peters, Grzegorz Pierczyński, Piotr Skowron, Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. 2023. Participatory budgeting: Data, tools, and analysis. *arXiv preprint arXiv:2305.11035* (2023).
- [12] Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. 2019. Fair knapsack. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 33. 1941–1948.
- [13] Duncan K Foley. 1970. Lindahl’s Solution and the Core of an Economy with Public Goods. *Econometrica: Journal of the Econometric Society* 38, 1 (1970), 66–72.
- [14] Rupert Freeman, David M Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2019. Truthful aggregation of budget proposals. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 751–752.
- [15] Daniel Gabay and Bertrand Mercier. 1976. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications* 2, 1 (1976), 17–40.
- [16] Donald Bruce Gillies. 1953. *Some theorems on n-person games*. Princeton University.
- [17] Roland Glowinski and Americo Marroco. 1975. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de Dirichlet non linéaires. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique* 9, R2 (1975), 41–76.
- [18] Ashish Goel, Anilesh K Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)* 7, 2 (2019), 1–27.
- [19] Bingsheng He and Xiaoming Yuan. 2012. On the $O(1/n)$ Convergence Rate of the Douglas-Rachford Alternating Direction Method. *SIAM J. Numer. Anal.* 50, 2 (2012), 700–709.
- [20] Magnus R Hestenes. 1969. Multiplier and gradient methods. *Journal of Optimization Theory and Applications* 4, 5 (1969), 303–320.
- [21] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. 2014. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*. 21–30.
- [22] Justin Hsu, Zhiyi Huang, Aaron Roth, and Zhiwei Steven Wu. 2016. Jointly private convex programming. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 580–599.
- [23] Zonghao Huang, Rui Hu, Yuanxiong Guo, Eric Chan-Tin, and Yanmin Gong. 2019. DP-ADMM: ADMM-based distributed learning with differential privacy. *IEEE Transactions on Information Forensics and Security* 15 (2019), 1002–1012.
- [24] Zhenqi Huang, Sayan Mitra, and Nitin Vaidya. 2015. Differentially private distributed optimization. In *Proceedings of the 16th International Conference on Distributed Computing and Networking*. 1–10.
- [25] Zhiyi Huang and Xue Zhu. 2018. Near optimal jointly private packing algorithms via dual multiplicative weight update. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 343–357.
- [26] Zhiyi Huang and Xue Zhu. 2019. Scalable and Jointly Differentially Private Packing. (2019).
- [27] Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. 2019. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 299–316.
- [28] László Lovász and Santosh Vempala. 2007. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms* 30, 3 (2007), 307–358.
- [29] Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 94–103.
- [30] Ilya Mironov. 2017. Rényi differential privacy. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF)*. 263–275.
- [31] Thomas J Muench. 1972. The core and the Lindahl equilibrium of an economy with a public good: An example. *Journal of Economic Theory* 4, 2 (1972), 241–255.
- [32] Kamesh Munagala, Yiheng Shen, Kangning Wang, and Zhiyi Wang. 2022. Approximate Core for Committee Selection via Multilinear Extension and Market Clearing. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2229–2252.
- [33] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. 2021. Proportional participatory budgeting with additive utilities. *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)* (2021), 12726–12737.
- [34] Michael JD Powell. 1969. A method for nonlinear constraints in minimization problems. *Optimization* (1969), 283–298.
- [35] Ariel D Procaccia. 2013. Cake cutting: Not just child’s play. *Commun. ACM* 56, 7 (2013), 78–87.
- [36] Ariel D Procaccia and Moshe Tennenholtz. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*. 177–186.
- [37] Herbert E Scarf. 1967. The core of an N person game. *Econometrica: Journal of the Econometric Society* (1967), 50–69.
- [38] Wei Shi, Qing Ling, Kun Yuan, Gang Wu, and Wotao Yin. 2014. On the linear convergence of the ADMM in decentralized consensus optimization. *IEEE Transactions on Signal Processing* 62, 7 (2014), 1750–1761.
- [39] Robert L Smith. 1984. Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research* 32, 6 (1984), 1296–1308.

- [40] Dariusz Stoliczki, Stanislaw Szufa, and Nimrod Talmon. 2020. Pabulib: A Participatory Budgeting Library. *arXiv preprint arXiv:2012.06539* (2020).
- [41] Roman Vershynin. 2018. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge University Press.
- [42] Tao Zhang and Quanyan Zhu. 2016. Dynamic differential privacy for ADMM-based distributed classification learning. *IEEE Transactions on Information Forensics and Security* 12, 1 (2016), 172–187.
- [43] Xueru Zhang, Mohammad Mahdi Khalili, and Mingyan Liu. 2018. Improving the privacy and accuracy of ADMM-based distributed algorithms. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. 5796–5805.

A NOTATIONS

Table . List of notations

| Notation | Description |
|----------------------------|---|
| n | Number of agents |
| m | Number of public items |
| s_j | Size of item j |
| s | Size vector, i.e., (s_1, \dots, s_m) |
| c | Total capacity |
| z_j | Fraction of item j that is allocated |
| z | Allocation variable, i.e., (z_1, \dots, z_m) |
| Z | Set of all feasible allocations, i.e., $\{z \in [0, 1]^m \mid s^T z \leq c\}$ |
| $U_i(z)$ | Agent i 's utility function for allocation z |
| u_i | Agent i 's utility vector, i.e., parameters of U_i : (u_{i1}, \dots, u_{id}) |
| U | Set $[0, 1]^d$ |
| u | Utility vectors for all agents, i.e., (u_1, \dots, u_n) |
| u_{-i} | Utility vector of all agents except agent i , i.e., $(u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$ |
| $M(u)$ | Randomized mechanism that maps $u \in U^n$ to probability distribution over Z |
| $\theta_i(z)$ | Logarithm of agent i 's utility, i.e., $\log(U_i(z))$ |
| θ | Summation of logarithm of all agents' utilities, i.e., $\sum_i \theta_i$ |
| L | Lipschitz parameter of $U_i(z)$'s |
| ϵ | Multiplicative approximation factor for truthfulness, core, and DP |
| δ | Additive approximation factor for truthfulness, core, and DP |
| α | Rényi divergence parameter |
| $\mathcal{N}(\mu, \Sigma)$ | Multivariate normal distribution with mean vector μ and covariance matrix Σ |
| K | Total number of iterations in Alg. 1 |
| $z^{(k)}$ | Global allocation variable at iteration k , i.e., $(z_{i1}^{(k)}, \dots, z_{im}^{(k)})$ |
| $x_i^{(k)}$ | Agent i 's local allocation variable at iteration k , i.e., $(x_{i1}^{(k)}, \dots, x_{im}^{(k)})$ |
| $x^{(k)}$ | Vector of local allocations at iteration k , i.e., $(x_1^{(k)}, \dots, x_n^{(k)})$ |
| $\gamma_i^{(k)}$ | Dual variable for $z = x_i$ constraint at iteration k , i.e., $(\gamma_{i1}^{(k)}, \dots, \gamma_{im}^{(k)})$ |
| $\gamma^{(k)}$ | Vector of dual variables, i.e., $(\gamma_1^{(k)}, \dots, \gamma_n^{(k)})$ |
| $q^{(k)}$ | Multivariate Gaussian noise added to $z^{(k)}$ at iteration k |
| σ^2 | Variance of added noise to each dimension of z |
| ρ | Penalty parameter for the augmented Lagrangian |
| L_i^ρ | Agent i 's partial augmented Lagrangian with parameter ρ |
| L^ρ | Summation of partial augmented Lagrangian functions, i.e., $\sum_i L_i^\rho$ |
| $\Pi_Z(z)$ | Euclidean projection of z onto Z , i.e., $\operatorname{argmin}_{z' \in Z} \ z - z'\ _2^2$ |
| \bar{z} | Time average of $z^{(k)}$'s, i.e., $(1/K) \sum_{k=1}^K z^{(k)}$ |
| \hat{z} | Euclidean projection of \bar{z} onto Z , i.e., $\Pi(\bar{z})$ |

B ELECTION INSTANCES

C OMITTED PROOFS

C.1 Proof of Lemma 3

PROOF. By concavity of U_i , for all $z, z' \in Z$, we have:

$$U_i(z') - U_i(z) \leq \nabla U_i(z)^T (z' - z). \quad (12)$$

Table . Characteristics of election instances.

| Inst. | Election | # Voters (n) | # Projects (m) | Budget (c) | Avg. # votes per voter |
|-------|----------------------------|-----------------|-------------------|---------------|---------------------------|
| 1 | Wroclaw'17 | 62,529 | 50 | 4,000,000 | 1.8 |
| 2 | Warszawa'20 Praga Poludnie | 14,897 | 134 | 5,900,907 | 9.1 |
| 3 | Katowice'21 | 36,370 | 47 | 3,003,438 | 1.5 |
| 4 | Warszawa'21 Mokotow | 12,933 | 98 | 7,147,577 | 9.7 |
| 5 | Wroclaw'16 Rejon NR 10-750 | 12,664 | 13 | 750,000 | 1 |
| 6 | Warszawa'23 Mokotow | 11,067 | 81 | 8,697,250 | 9 |
| 7 | Wroclaw'16 Rejon NR 12-250 | 10,711 | 15 | 650,000 | 1 |
| 8 | Wroclaw'16 | 67,103 | 52 | 4,500,000 | 1.8 |
| 9 | Warszawa'22 | 81,234 | 129 | 28,072,528 | 7.9 |
| 10 | Gdansk'20 | 30,237 | 28 | 3,600,000 | 1 |
| 11 | Warszawa'21 | 95,899 | 106 | 24,933,409 | 8.3 |
| 12 | Warszawa'20 | 86,721 | 101 | 24,933,409 | 7.2 |

Let z^* be an MNW solution. The condition of optimality requires that for any $z' \in Z$, we have:

$$\sum_i \frac{\nabla U_i(z^*)^T}{U_i(z^*)} (z' - z^*) \leq 0 \xrightarrow{\text{by (12)}} \frac{1}{n} \sum_i \frac{U_i(z')}{U_i(z^*)} \leq 1. \quad (13)$$

For contradiction, suppose that z^* is not a core outcome. Then there exists a set of agents A and an allocation z' such that $(|A|/n)U_i(z') \geq U_i(z^*)$, and at least one inequality is tight. This implies $(1/n) \sum_{i \in A} U_i(z')/U_i(z^*) > 1$, which contradicts (13). \square

C.2 Proof of Lemma 16

We start by establishing some key notations. We define $\tilde{y}^{(k)}$ as:

$$\tilde{y}^{(k)} = y^{(k-1)} + \rho(x^{(k)} - Gz^{(k-1)}),$$

where $G = (I_m, \dots, I_m)$. Additionally, $\tilde{w}^{(k)}$ and $F(w)$ as follows:

$$\tilde{w}^{(k)} = \begin{pmatrix} x^{(k)} \\ z^{(k)} \\ \tilde{y}^{(k)} \end{pmatrix}, \text{ and } F(w) = \begin{pmatrix} -y \\ \sum_i y_i \\ x - G_{n,m}z \end{pmatrix}. \quad (14)$$

To prove Lemma 16, we first provide an important inequality that relates $\tilde{w}^{(k)}$ to any $w \in W$:

LEMMA 19. *Let $\{\tilde{w}^{(k)}\}$ and $\{q^{(k)}\}$ be sequences produced by Alg. 1. For all $w \in W$, we have:*

$$\begin{aligned} (x - x^{(k)})^T \nabla \theta(x^{(k)}) + w^T F(\tilde{w}^{(k)}) &\leq n\rho(z^{(k)} - z)^T q^{(k)} - \frac{1}{2\rho} \|\tilde{y}^{(k)} - y^{(k-1)}\|_2^2 \\ &\quad + \frac{n\rho}{2} \left(\|z - z^{(k-1)}\|_2^2 - \|z - z^{(k)}\|_2^2 \right) \\ &\quad + \frac{1}{2\rho} \left(\|y - y^{(k-1)}\|_2^2 - \|y - y^{(k)}\|_2^2 \right). \end{aligned} \quad (15)$$

PROOF. For all i , the condition of optimality for Line 6 of Alg. 1 requires:

$$(x_i - x_i^{(k)})^T (\nabla \theta_i(x_i^{(k)}) - y_i^{(k-1)} - \rho(x_i^{(k)} - z^{(k-1)})) \leq 0 \quad \text{for all } x_i \in Z. \quad (16)$$

Given the definition of $\tilde{y}^{(k)}$, we can rewrite (16) for each i and $x_i \in Z$ as:

$$(x_i - x_i^{(k)})^T (\nabla \theta_i(x_i^{(k)}) - \tilde{y}_i^{(k)}) \leq 0.$$

Summing over all i , for any $x \in Z^n$, we have:

$$(x - x^{(k)})^T \nabla \theta(x^{(k)}) - (x - x^{(k)})^T \tilde{y}^{(k)} \leq 0. \quad (17)$$

Next, given (7), Line 8 of Alg. 1 implies that $z^{(k)}$ is the solution to:

$$\underset{z}{\text{maximize}} \quad \sum_i \left(-(\gamma_i^{(k-1)})^T (x_i^{(k)} - z + q^{(k)}) - \frac{\rho}{2} \|x_i^{(k)} - z + q^{(k)}\|_2^2 \right).$$

The condition of optimality for this optimization requires that:

$$(z - z^{(k)})^T \left(\sum_i \left(\gamma_i^{(k-1)} + \rho(x_i^{(k)} - z^{(k)} + q^{(k)}) \right) \right) \leq 0 \quad \text{for all } z \in \mathbb{R}^m. \quad (18)$$

Given the definition of $\tilde{y}^{(k)}$, we can rewrite (18) for all $z \in \mathbb{R}^m$ as:

$$\begin{aligned} (z - z^{(k)})^T \left(\sum_i \tilde{y}_i^{(k)} - n\rho(z^{(k)} - z^{(k-1)}) + n\rho q^{(k)} \right) &\leq 0 \Rightarrow \\ (z - z^{(k)})^T \sum_i \tilde{y}_i^{(k)} &\leq n\rho(z - z^{(k)})^T (z^{(k)} - z^{(k-1)}) - n\rho(z - z^{(k)})^T q^{(k)}. \end{aligned} \quad (19)$$

Next, given Line 9 of Alg. 1, for all $\gamma \in \mathbb{R}^{mn}$ we have:

$$\begin{aligned} x^{(k)} - Gz^{(k)} &= (\gamma^{(k)} - \gamma^{(k-1)})/\rho \Rightarrow \\ (\gamma - \tilde{\gamma}^{(k)})^T (x^{(k)} - Gz^{(k)}) &= (\gamma - \tilde{\gamma}^{(k)})^T (\gamma^{(k)} - \gamma^{(k-1)})/\rho. \end{aligned} \quad (20)$$

Combining (17)–(20), for any $w = (x, z, \gamma) \in W$, we have:

$$\begin{aligned} (x - x^{(k)})^T \nabla \theta(x^{(k)}) + (w - \tilde{w}^{(k)})^T F(\tilde{w}^{(k)}) &\leq n\rho(z^{(k)} - z)^T q^{(k)} \\ &+ n\rho(z - z^{(k)})^T (z^{(k)} - z^{(k-1)}) + (\gamma - \tilde{\gamma}^{(k)})^T (\gamma^{(k)} - \gamma^{(k-1)})/\rho. \end{aligned} \quad (21)$$

Applying the following identity:

$$w^T F(w) = -x^T \gamma + z^T \sum_i \gamma_i + \gamma^T (x - Gz) = 0, \quad (22)$$

we can rewrite (21) as:

$$\begin{aligned} (x - x^{(k)})^T \nabla \theta(x^{(k)}) + w^T F(\tilde{w}^{(k)}) &\leq n\rho(z^{(k)} - z)^T q^{(k)} \\ &+ n\rho(z - z^{(k)})^T (z^{(k)} - z^{(k-1)}) + (\gamma - \tilde{\gamma}^{(k)})^T (\gamma^{(k)} - \gamma^{(k-1)})/\rho. \end{aligned} \quad (23)$$

Next, we focus on the right-hand side of (23). Given the following identity:

$$(a - b)^T (c - d) = \frac{1}{2} (\|a - d\|_2^2 - \|a - c\|_2^2) + \frac{1}{2} (\|b - c\|_2^2 - \|b - d\|_2^2),$$

we have:

$$n\rho(z - z^{(k)})^T (z^{(k)} - z^{(k-1)}) = \frac{n\rho}{2} (\|z - z^{(k-1)}\|_2^2 - \|z - z^{(k)}\|_2^2) - \frac{n\rho}{2} \|z^{(k)} - z^{(k-1)}\|_2^2, \quad (24)$$

and:

$$\begin{aligned} \frac{1}{\rho} (\gamma - \tilde{\gamma}^{(k)})^T (\gamma^{(k)} - \gamma^{(k-1)}) &= \frac{1}{2\rho} (\|\gamma - \gamma^{(k-1)}\|_2^2 - \|\gamma - \gamma^{(k)}\|_2^2) \\ &+ \frac{1}{2\rho} (\|\tilde{\gamma}^{(k)} - \gamma^{(k)}\|_2^2 - \|\tilde{\gamma}^{(k)} - \gamma^{(k-1)}\|_2^2). \end{aligned} \quad (25)$$

Given the definition of $\tilde{Y}^{(k)}$ and Line 9 of Alg. 1, we have:

$$\begin{aligned}\|\tilde{Y}^{(k)} - Y^{(k)}\|_2^2 &= \|\rho(x^{(k)} - Gz^{(k-1)}) - (Y^{(k)} - Y^{(k-1)})\|_2^2 \\ &= \rho^2 \|x^{(k)} - Gz^{(k-1)} - x^{(k)} + Gz^{(k)}\|_2^2 \\ &= n\rho^2 \|z^{(k)} - z^{(k-1)}\|_2^2.\end{aligned}\quad (26)$$

Combining (24)–(26), we obtain

$$\begin{aligned}&n\rho(z - z^{(k)})^T(z^{(k)} - z^{(k-1)}) + (Y - \tilde{Y}^{(k)})^T(Y^{(k)} - Y^{(k-1)})/\rho = \\ &\frac{n\rho}{2} \left(\|z - z^{(k-1)}\|_2^2 - \|z - z^{(k)}\|_2^2 \right) + \frac{1}{2\rho} \left(\|Y - Y^{(k-1)}\|_2^2 - \|Y - Y^{(k)}\|_2^2 \right) - \frac{1}{2\rho} \|\tilde{Y}^{(k)} - Y^{(k-1)}\|_2^2.\end{aligned}\quad (27)$$

Substituting (27) into (23) gives (15). \square

We are now ready to prove Lemma 16:

PROOF. We start by rewriting $(x - x^{(k)})^T \nabla \theta(x^{(k)})$ as:

$$(x - x^{(k)})^T \nabla \theta(x^{(k)}) = \sum_i (x_i - x_i^{(k)})^T \nabla U_i(x_i^{(k)})/U_i(x_i^{(k)}).$$

Since $U_i(x)$ is concave, for any i and for any $x, x' \in Z$, we have:

$$U_i(x') - U_i(x) \leq (x' - x)^T \nabla U_i(x). \quad (28)$$

Therefore, (15) implies:

$$\begin{aligned}\sum_i \frac{U_i(x_i)}{U_i(x_i^{(k)})} + w^T F(\tilde{w}^{(k)}) &\leq n + n\rho(z^{(k)} - z)^T q^{(k)} + \frac{n\rho}{2} \left(\|z - z^{(k-1)}\|_2^2 - \|z - z^{(k)}\|_2^2 \right) \\ &\quad + \frac{1}{2\rho} \left(\|Y - Y^{(k-1)}\|_2^2 - \|Y - Y^{(k)}\|_2^2 \right).\end{aligned}\quad (29)$$

Next, we define $W_Z = \{(Gz, z, \mathbf{0}_{mn}) \mid z \in Z\}$. For any $w \in W_Z$, we have:

$$w^T F(\tilde{w}^{(k)}) = -z^T \sum_i \tilde{Y}_i^{(k)} + z^T \sum_i \tilde{Y}_i^{(k)} = 0.$$

Given this identity, for any $w \in W_Z \subset W$, (29) implies:

$$\begin{aligned}\frac{1}{n} \sum_i \frac{U_i(z)}{U_i(x_i^{(k)})} &\leq 1 + \rho(z^{(k)} - z)^T q^{(k)} + \frac{\rho}{2} \left(\|z - z^{(k-1)}\|_2^2 - \|z - z^{(k)}\|_2^2 \right) \\ &\quad + \frac{1}{2n\rho} \left(\|Y^{(k-1)}\|_2^2 - \|Y^{(k)}\|_2^2 \right).\end{aligned}$$

Summing this inequality over $k = 1$ to K and dividing by K , for any $z \in Z$, we obtain:

$$\frac{1}{n} \sum_i \frac{1}{K} \sum_{k=1}^K \frac{U_i(z)}{U_i(x_i^{(k)})} \leq 1 + \frac{\rho}{K} \sum_{k=1}^K (z^{(k)} - z)^T q^{(k)} + \frac{\rho}{2K} \|z\|_2^2. \quad (30)$$

Since $U_i(x)$ is a strictly increasing concave function, $1/U_i(x)$ is convex [2]. As a result, by Jensen's inequality, for any $z \in Z$, we have:

$$\frac{1}{K} \sum_{k=1}^K \frac{U_i(z)}{U_i(x_i^{(k)})} \geq \frac{U_i(z)}{U_i\left(\frac{1}{K} \sum_{k=1}^K x_i^{(k)}\right)} = \frac{U_i(z)}{U_i(\bar{x})}.$$

Given this inequality and the fact that for any $z \in Z$, $\|z\|_2^2 \leq 1$, (30) implies (10). \square

C.3 Proof of Lemma 17

To prove Lemma 17, we first show that the sequence $\{\tilde{\mathbf{w}}^{(k)}\}$, define in (14), is *contractive* (with some noise):

LEMMA 20. *Let $\{\mathbf{w}^{(k)}\}$ and $\{q^{(k)}\}$ be sequences generated by Alg. 1. Let $\mathbf{w}^* = (x^*, z^*, \gamma^*)$ be the solution to (4), with $x_i^* = z^*$ for all i . Then we have:*

$$n\rho^2 \left(\|z^{(k)} - z^*\|_2^2 - \|z^{(k-1)} - z^*\|_2^2 \right) + \left(\|\gamma^{(k)} - \gamma^*\|_2^2 - \|\gamma^{(k-1)} - \gamma^*\|_2^2 \right) \leq 2n\rho^2 (z^{(k)} - z^*)^T q^{(k)}. \quad (31)$$

PROOF. Setting $\mathbf{w} := \mathbf{w}^*$ in (15), we have:

$$\begin{aligned} (x^* - x^{(k)})^T \nabla \theta(x^{(k)}) + \mathbf{w}^{*T} F(\tilde{\mathbf{w}}^{(k)}) &\leq n\rho(z^{(k)} - z^*)^T q^{(k)} - \frac{1}{2\rho} \|\tilde{\gamma}^{(k)} - \gamma^{(k-1)}\|_2^2 \\ &\quad + \frac{n\rho}{2} \left(\|z^* - z^{(k-1)}\|_2^2 - \|z^* - z^{(k)}\|_2^2 \right) + \frac{1}{2\rho} \left(\|\gamma^* - \gamma^{(k-1)}\|_2^2 - \|\gamma^* - \gamma^{(k)}\|_2^2 \right). \end{aligned} \quad (32)$$

Since \mathbf{w}^* is the solution to (4), the conditions of optimality require:

$$(x^{(k)} - x^*)^T \nabla \theta(x^*) + (\tilde{\mathbf{w}}^{(k)} - \mathbf{w}^*)^T F(\mathbf{w}^*) \leq 0. \quad (33)$$

It can be easily shown that $(\tilde{\mathbf{w}}^{(k)} - \mathbf{w}^*)^T F(\mathbf{w}^*) + \mathbf{w}^{*T} F(\tilde{\mathbf{w}}^{(k)}) = 0$. Therefore, summing (32) and (33), we have:

$$\begin{aligned} (x^* - x^{(k)})^T (\nabla \theta(x^{(k)}) - \nabla \theta(x^*)) &\leq n\rho(z^{(k)} - z^*)^T q^{(k)} - \frac{1}{2\rho} \|\tilde{\gamma}^{(k)} - \gamma^{(k-1)}\|_2^2 \\ &\quad + \frac{n\rho}{2} \left(\|z^* - z^{(k-1)}\|_2^2 - \|z^* - z^{(k)}\|_2^2 \right) + \frac{1}{2\rho} \left(\|\gamma^* - \gamma^{(k-1)}\|_2^2 - \|\gamma^* - \gamma^{(k)}\|_2^2 \right). \end{aligned} \quad (34)$$

Since $\theta(x)$ is concave, we have $(x^* - x^{(k)})^T (\nabla \theta(x^*) - \nabla \theta(x^{(k)})) \leq 0$. Given this, (34) implies (31). \square

We are now ready to prove Lemma 17:

PROOF. Summing (31) over $k = 1$ to K , we have:

$$n\rho^2 \left(\|z^{(K)} - z^*\|_2^2 - \|z^*\|_2^2 \right) + \left(\|\gamma^{(K)} - \gamma^*\|_2^2 - \|\gamma^*\|_2^2 \right) \leq 2n\rho^2 \sum_{k=1}^K (z^{(k)} - z^*)^T q^{(k)}.$$

This inequality implies:

$$\|\gamma^{(K)} - \gamma^*\|_2^2 \leq 2n\rho^2 \sum_{k=1}^K (z^{(k)} - z^*)^T q^{(k)} + n\rho^2 \|z^*\|_2^2 + \|\gamma^*\|_2^2. \quad (35)$$

Next, we have:

$$\begin{aligned} \rho K \|\bar{x} - G\bar{z}\|_2 &= \|\gamma^{(K)}\|_2 = \|\gamma^{(K)} - \gamma^* + \gamma^*\|_2 \\ &\leq \|\gamma^*\|_2 + \|\gamma^{(K)} - \gamma^*\|_2 \\ &\leq \|\gamma^*\|_2 + \left(2n\rho^2 \sum_{k=1}^K (z^{(k)} - z^*)^T q^{(k)} + n\rho^2 \|z^*\|_2^2 + \|\gamma^*\|_2^2 \right)^{\frac{1}{2}}. \end{aligned} \quad (36)$$

The Euclidean projection onto the convex set Z is contractive. Therefore, since $\bar{x}, \hat{z} \in Z$, we have:

$$\|\bar{x} - G\hat{z}\|_2 = \|\bar{x} - G\Pi_Z(\bar{z})\|_2 \leq \|\bar{x} - G\bar{z}\|_2.$$

Given this inequality, (36) implies (11). \square

C.4 Proof of Thm. 18

PROOF. Let $\{w^{(k)}\}$ and $\{q^{(k)}\}$ be sequences generated by Alg. 1, and let $w^* = (x^*, z^*, y^*)$ be the solution to (4). Define y_{\max}^* and \hat{z} as:

$$y_{\max}^* = \max_{i,j} |y_{i,j}^*| \quad \text{and} \quad \hat{z} = \operatorname{argmax}_{z \in Z} \sum_i \frac{U_i(z)}{U_i(\bar{x}_i)}.$$

Further, define ε_1 and ε_2 as:

$$\varepsilon_1 = \frac{\rho}{K} \sum_{k=1}^K (z^{(k)} - \hat{z})^T q^{(k)} + \frac{\rho}{2K},$$

and

$$\varepsilon_2 = \frac{L}{\rho K} \sqrt{nm} y_{\max}^* + \frac{L}{\rho K} \left(2n\rho^2 \sum_{k=1}^K (z^{(k)} - z^*)^T q^{(k)} + n\rho^2 + nm y_{\max}^*{}^2 \right)^{1/2}.$$

By Lemma 16 and the definition of \hat{z} , we have:

$$\frac{1}{n} \sum_i \frac{U_i(z)}{U_i(\bar{x}_i)} \leq \frac{1}{n} \sum_i \frac{U_i(\hat{z})}{U_i(\bar{x}_i)} \leq 1 + \varepsilon_1. \quad (37)$$

Due to the Lipschitz continuity of $U_i(x)$ and the fact that $\|\bar{x}_i - \hat{z}\|_2 \leq \|\bar{x} - G\hat{z}\|_2$ for any i , Lemma 17 implies:

$$U_i(\bar{x}) \leq U_i(\hat{z}) + \varepsilon_2 \quad \forall i. \quad (38)$$

Combining (37) and (38), we obtain:

$$\frac{1}{n} \sum_i \frac{U_i(z)}{U_i(\hat{z}) + \varepsilon_2} \leq 1 + \varepsilon_1. \quad (39)$$

Given (39), Lemma 5 implies that \hat{z} is an $(\varepsilon_1, \varepsilon_2 + \varepsilon_1 \varepsilon_2)$ -core outcome. However, ε_1 and ε_2 are random variables. We next study the tail behavior of these random variables by establishing a concentration bound on $\sum_{k=1}^K (z^{(k)} - z)^T q^{(k)}$ for any $z \in Z$.

$$\begin{aligned} |(z^{(k)} - z)^T q^{(k)}| &= \left| \left(\frac{1}{n} \sum_i x_i^{(k)} + q^{(k)} - q^{(k-1)} - z \right)^T q^{(k)} \right| \\ &\leq \left| \left(\frac{1}{n} \sum_i x_i^{(k)} \right)^T q^{(k)} \right| + \|q^{(k)}\|_2^2 + |q^{(k-1)T} q^{(k)}| + |z^T q^{(k)}| \\ &\leq \left\| \frac{1}{n} \sum_i x_i^{(k)} \right\|_1 \|q^{(k)}\|_2 + \|q^{(k)}\|_2^2 + |q^{(k-1)T} q^{(k)}| + \|z\|_1 \|q^{(k)}\|_2 \\ &\leq 2\|q^{(k)}\|_2 + \|q^{(k)}\|_2^2 + |q^{(k-1)T} q^{(k)}|, \end{aligned} \quad (40)$$

where the first inequality follows from the triangle inequality, the second inequality follows from Cauchy-Schwarz inequality and the fact that $\|\cdot\|_2 \leq \|\cdot\|_1$ (i.e., $|a^T b| \leq \|a\|_2 \|b\|_2 \leq \|a\|_1 \|b\|_2$ for all vectors a and b of an inner product space), the third inequality follows from the fact that

$\frac{1}{n} \sum_i x_i^{(k)} \in Z$, and for any $z \in Z$, we have $\|z\|_2 \leq \|z\|_1 \leq 1$. For the last term in (40), we have:

$$\begin{aligned} |q^{(k-1)T} q^{(k)}| &= \left| \sum_{j=1}^m q_j^{(k-1)} q_j^{(k)} \right| \leq \sum_{j=1}^m |q_j^{(k-1)} q_j^{(k)}| \\ &= \sum_{j=1}^m |q_j^{(k-1)}| |q_j^{(k)}| \\ &\leq \frac{1}{2} \sum_{j=1}^m \left((q_j^{(k-1)})^2 + (q_j^{(k)})^2 \right) \\ &= \frac{1}{2} \|q^{(k-1)}\|_2^2 + \frac{1}{2} \|q^{(k)}\|_2^2, \end{aligned}$$

where the first inequality follows from the triangle inequality, and the second inequality follows from the Young's inequality. Substituting the last inequality into (40) and summing over k , we have:

$$\begin{aligned} \left| \sum_{k=1}^K (z^{(k)} - z)^T q^{(k)} \right| &\leq \sum_{k=1}^K |(z^{(k)} - z)^T q^{(k)}| \\ &\leq 2 \sum_{k=1}^K \left(\|q^{(k)}\|_2^2 + \|q^{(k)}\|_2 \right). \end{aligned} \quad (41)$$

Since $q_j^{(k)} \sim \mathcal{N}(0, \sigma^2)$ is a sub-Gaussian random variable, we have that $(q_j^{(k)})^2$ is a sub-exponential random variable ([41], Lemma 2.7.6) with $\mathbb{E}[(q_j^{(k)})^2] = \sigma^2$ and

$$\|(q_j^{(k)})^2 - \sigma^2\|_{\psi_1} \leq C_1 \sigma^2,$$

where C_1 is a constant, and $\|X\|_{\psi_1} = \inf\{t > 0 \mid \mathbb{E}[\exp(|X|/t)] \leq 2\}$ is the *sub-exponential norm* of a real-valued random variable X . Since $q_j^{(k)}$'s are i.i.d. for all k 's and j 's, by the Bernstein's inequality ([41], Theorem 2.8.1), for any $t \geq 0$, we have:

$$\mathbb{P} \left[\sum_{k=1}^K \|q^{(k)}\|_2^2 - Km\sigma^2 \geq t \right] \leq \exp \left(-c_1 \min \left(\frac{t^2}{Km\sigma^4}, \frac{t}{\sigma^2} \right) \right), \quad (42)$$

where c_1 is a constant.

Next, $\|q^{(k)}\|_2$ is a sub-Gaussian random variable ([41], Theorem 3.1.1 and Lemma 2.6.8) with

$$\left\| \|q^{(k)}\|_2 - \mathbb{E}[\|q^{(k)}\|_2] \right\|_{\psi_2} \leq C_2 \sigma^2,$$

where C_2 is a constant, and $\|X\|_{\psi_2} = \inf\{t > 0 \mid \mathbb{E}[\exp(X^2/t^2)] \leq 2\}$ is the *sub-Gaussian norm* of a real-valued random variable X . Since $q^{(k)}$'s are independent, by the general Hoeffding's inequality ([41], Theorem 2.6.2), for any $t \geq 0$, we have:

$$\mathbb{P} \left[\sum_{k=1}^K \left(\|q^{(k)}\|_2 - \mathbb{E}[\|q^{(k)}\|_2] \right) \geq t \right] \leq \exp \left(-\frac{c_2 t^2}{K\sigma^4} \right),$$

where c_2 is a constant. We next provide an upper bound on $\mathbb{E}[\|q^{(k)}\|_2]$. Consider the following inequality which holds for any $u \geq 0$:

$$\sqrt{u} \leq \frac{1+u}{2}.$$

By setting $u = \frac{1}{m\sigma^2} \|q^{(k)}\|_2^2$, we get:

$$\frac{\|q^{(k)}\|_2}{\sqrt{m}\sigma} \leq \frac{1 + (1/m\sigma^2)\|q^{(k)}\|_2^2}{2}.$$

Taking the expectation of both sides of the last inequality gives:

$$\mathbb{E} \left[\|q^{(k)}\|_2 \right] \leq \sqrt{m}\sigma \frac{1+1}{2} = \sqrt{m}\sigma.$$

Therefore, we have:

$$\mathbb{P} \left[\sum_{k=1}^K \|q^{(k)}\|_2 - K\sqrt{m}\sigma \geq t \right] \leq \mathbb{P} \left[\sum_{k=1}^K \left(\|q^{(k)}\|_2 - \mathbb{E} \left[\|q^{(k)}\|_2 \right] \right) \geq t \right] \leq \exp\left(-\frac{c_2 t^2}{K\sigma^4}\right). \quad (43)$$

Given (41)–(43) and the union bound, for $t' = 4t + 2Km\sigma^2 + 2K\sqrt{m}\sigma$, we have:

$$\begin{aligned} \mathbb{P} \left[\left| \sum_{k=1}^K (z^{(k)} - z)^T q^{(k)} \right| \geq t' \right] &\leq \mathbb{P} \left[2 \sum_{k=1}^K \left(\|q^{(k)}\|_2^2 + \|q^{(k)}\|_2 \right) \geq t' \right] \\ &\leq \mathbb{P} \left[\sum_{k=1}^K \|q^{(k)}\|_2^2 - Km\sigma^2 \geq t \right] + \mathbb{P} \left[\sum_{k=1}^K \|q^{(k)}\|_2 - K\sqrt{m}\sigma \geq t \right] \\ &\leq \exp \left(-c_1 \min \left(\frac{t^2}{Km\sigma^4}, \frac{t}{\sigma^2} \right) \right) + \exp\left(-\frac{c_2 t^2}{K\sigma^4}\right). \end{aligned}$$

For some constant c , setting $t = c\sqrt{Km}\sigma^2 \log^{1/2}(n)$ in the last inequality implies that with probability more than $1 - \frac{1}{n} - \frac{1}{n^m}$, the following inequality holds for any $z \in Z$:

$$\sum_{k=1}^K (z^{(k)} - z)^T q^{(k)} \leq 4c\sqrt{Km}\sigma^2 \log^{1/2}(n) + 2Km\sigma^2 + 2K\sqrt{m}\sigma. \quad (44)$$

Given that $K = \Theta(n)$, $\epsilon = \Theta(1/\log(n))$, $\delta = \Theta(1/\sqrt{n})$, $\alpha = \Theta(\log^2(n))$, and $m = o(\sqrt{n})$, substituting (44) into ε_1 and ε_2 implies that ε_1 and ε_2 asymptotically go to zero as n grows. Therefore, given (39), by Lemma 5, the output of Alg. 1 is an asymptotic core solution with probability more than $1 - \frac{1}{n} - \frac{1}{n^m}$.

□