

Fair Allocation of Indivisible Public Goods

BRANDON FAIN, Duke University, USA

KAMESH MUNAGALA, Duke University, USA

NISARG SHAH, University of Toronto, Canada

We consider the problem of fairly allocating indivisible public goods. We model the public goods as elements with feasibility constraints on what subsets of elements can be chosen, and assume that agents have additive utilities across elements. Our model generalizes existing frameworks such as fair public decision making and participatory budgeting. We study a groupwise fairness notion called the *core*, which generalizes well-studied notions of proportionality and Pareto efficiency, and requires that each subset of agents must receive an outcome that is fair relative to its size.

In contrast to the case of divisible public goods (where fractional allocations are permitted), the core is not guaranteed to exist when allocating indivisible public goods. Our primary contributions are the notion of an *additive* approximation to the core (with a tiny multiplicative loss), and polynomial time algorithms that achieve a small additive approximation, where the additive factor is relative to the largest utility of an agent for an element. If the feasibility constraints define a matroid, we show an additive approximation of 2. A similar approach yields a constant additive bound when the feasibility constraints define a matching. For feasibility constraints defining an arbitrary packing polytope with mild restrictions, we show an additive guarantee that is logarithmic in the *width* of the polytope. Our algorithms are based on the convex program for maximizing the Nash social welfare, but differ significantly from previous work in how it is used. As far as we are aware, our work is the first to approximate the core in indivisible settings.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**;

Additional Key Words and Phrases: Fairness, public goods, indivisible goods, core, proportionality

1 INTRODUCTION

In fair resource allocation, most work considers *private goods*; each good must be assigned to a particular agent (and no other). However, not all goods are private. *Public goods* are those which can be enjoyed by multiple agents simultaneously, like a public road. Allocation of public goods generalizes the problem of allocation of private goods, and, as we will see, can provide new difficulties from both a normative and an algorithmic perspective.

Consider an example to highlight what a public resource allocation problem might look like, and why fairness might be a concern. Suppose that the next time you vote, you see that there are four referendums for your consideration on the ballot, all of which concern the allocation of various public goods in your city: A = a new school, B = enlarging the public library, C = renovating the community college, and D = improving a museum. In 2016, residents of Durham, North Carolina faced precisely these options [24]. Suppose the government has resources to fund only two of the four projects, and the (hypothetical) results were as follows: a little more than half of the population voted for (A, B) , a little less than half voted for (C, D) , and every other combination received a small number of votes. Which projects should be funded?

The full version of this paper is available at <https://arxiv.org/abs/1805.03164>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06... \$15.00

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

<https://doi.org/10.1145/3219166.3219174>

If we naïvely tally the votes, we would fund A and B, and ignore the preferences of a very large minority. In contrast, funding A and C seems like a reasonable compromise. Of course, it is impossible to satisfy *all* voters, but given a wide enough range of possible outcomes, perhaps we can find one that fairly reflects the preferences of large subsets of the population. This idea is not captured by fairness axioms like proportionality or their approximations [9], which view fairness from the perspectives of *individual* agents. Indeed, in the aforementioned example, *every* allocation gives zero utility to *some* agent, and would be deemed equally good according to such fairness criteria.

1.1 Public Goods Model

We consider a fairly broad model for public goods allocation that generalizes much of previous work [3, 8–10, 12, 23, 29]. There is a set of voters (or agents) $N = [n]$. Public goods are modeled as elements of a ground set W . We denote $m = |W|$. An *outcome* c is a subset of W . Let $\mathcal{F} \subseteq 2^W$ denote the set of feasible outcomes.

The utility of agent i for element $j \in W$ is denoted $u_{ij} \in \mathbb{R}_{\geq 0}$. We assume that agents have additive utilities, i.e., the utility of agent i under outcome $c \in \mathcal{F}$ is $u_i(c) = \sum_{j \in c} u_{ij}$. Since we are interested in scale-invariant guarantees, we assume without loss of generality that $\max_{j \in W} u_{ij} = 1$ for each agent i , so that $u_{ij} \in [0, 1]$ for all i, j . Crucially, this does not restrict the utility of an agent for an outcome to be 1: $u_i(c)$ can be as large as m . Specifically, let $V_i = \max_{c \in \mathcal{F}} u_i(c)$, and $V_{\max} = \max_{i \in N} V_i$. Our results differ by the feasibility constraints imposed on the outcome. We consider three types of constraints, special cases of which have been studied previously in literature.

Matroid Constraints. In this setting, we are given a matroid \mathcal{M} over the ground set W , and the feasibility constraint is that the chosen elements must form a basis of \mathcal{M} (see Kung [21] for a formal introduction to matroids).

This generalizes the *public decision making* setting introduced by Conitzer et al. [9]. In this setting, there is a set of issues T , and each issue $t \in T$ has an associated set of alternatives $A^t = \{a_1^t, \dots, a_{k_t}^t\}$, exactly one of which must be chosen. Agent i has utility $u_i^t(a_j^t)$ if alternative a_j^t is chosen for issue t , and utilities are additive across issues. An outcome c chooses one alternative for every issue. It is easy to see that if the ground set is $\cup_t A^t$, the feasibility constraints correspond to a partition matroid. We note that public decision making generalizes *private goods allocation* with additive utility [8, 23, 29] in which each private good must be allocated to exactly one agent.

Matroid constraints also capture multi-winner elections in the voting literature (see, e.g. [3]), in which voters have additive utilities over candidates, and a committee of at most k candidates must be chosen. This is captured by a uniform matroid over the set of candidates.

Matching Constraints. In this setting, the elements are edges of an undirected graph $G(V, E)$, and the feasibility constraint is that the subset of edges chosen must form a *matching*. Matchings constraints in a bipartite graph can be seen as the intersection of two matroid constraints. Matching constraints are a special case of the more general packing constraints we consider below.

Packing Constraints. In this setting, we impose a set of packing constraints $Ax \leq \mathbf{b}$, where $x_j \in \{0, 1\}$ is the indicator denoting whether element j is chosen in the outcome. Suppose A is a $K \times m$ matrix, so that there are K packing constraints. By scaling, we can assume $a_{kj} \in [0, 1]$ for all k, j . Note that even for one agent, packing constraints encode independent set. Thus, to make the problem tractable, we assume \mathbf{b} is sufficiently large, in particular, $b_k = \omega(\log K)$ for all $k \in \{1, 2, \dots, K\}$. This is in contrast to matroid and matching constraints, for which single-agent problems are polynomial time solvable. A classic measure of how easy it is to satisfy the packing

constraints is the *width* ρ [30]:

$$\rho = \max_{k \in [K]} \frac{\sum_{j \in [m]} a_{kj}}{b_k}. \quad (1)$$

Packing constraints capture the general KNAPSACK setting, in which there is a set of m items, each item j has an associated size s_j , and a set of items of total size at most B must be selected. This setting is motivated by participatory budgeting applications [6, 10, 12, 14, 16, 17, 28], in which the items are public projects, and the sizes represents the costs of the projects. KNAPSACK uses a single packing constraint. Multiple packing constraints can arise if the projects consume several types of resources, and there is a budget constraint for each resource type. For example, consider a statewide participatory budgeting scenario where each county has a budget than can only be spent on projects affecting that county, the state has some budget that can be spent in any county, and projects might affect multiple counties. In such settings, it is natural to assume a small width, i.e., that the budget for each resource is such that a large fraction (but not all) of the projects can be funded. We note that the aforementioned multi-winner election problem is a special case of the KNAPSACK problem with unit sizes.

1.2 Prior Work: Fairness Properties

We define fairness desiderata for the public goods setting by generalizing appropriate desiderata from the private goods setting such as Pareto optimality, which is a weak notion of efficiency, and proportionality, which is a per-agent fair share guarantee.¹

Definition 1.1. An outcome \mathbf{c} satisfies **Pareto optimality** if there is no other outcome \mathbf{c}' such that $u_i(\mathbf{c}') \geq u_i(\mathbf{c})$ for all agents $i \in N$, and at least one inequality is strict.

Recall that V_i is the maximum possible utility agent i can derive from a feasible outcome.

Definition 1.2. The **proportional share** of an agent $i \in N$ is $Prop_i := \frac{V_i}{n}$. For $\beta \in (0, 1]$, we say that an outcome \mathbf{c} satisfies β -proportionality if $u_i(\mathbf{c}) \geq \beta \cdot Prop_i$ for all agents $i \in N$. If $\beta = 1$, we simply say that \mathbf{c} satisfies proportionality.

The difficulty in our setting stems from requiring integral outcomes, and not allowing randomization. In the absence of randomization, it is reasonably straightforward to show that we cannot guarantee β -proportionality for any $\beta \in (0, 1]$. Consider a problem instance with two agents and two feasible outcomes, where each outcome gives a positive utility to a unique agent. In any feasible outcome, one agent has zero utility, which violates β -proportionality for every $\beta > 0$.

To address this issue, Conitzer et al. [9] introduced the novel relaxation of proportionality up to one issue in their public decision making framework, inspired by a similar relaxation called envy-freeness up to one good in the private goods setting [8, 23]. They say that an outcome \mathbf{c} of a public decision making problem satisfies *proportionality up to one issue* if for all agents $i \in N$, there exists an outcome \mathbf{c}' that differs from \mathbf{c} only on a single issue and $u_i(\mathbf{c}') \geq Prop_i$. Proportionality up to one issue is a reasonable fairness guarantees only when the number of issues is larger than the number of agents; otherwise, it is vacuous and is satisfied by all outcomes. Thus, it is perfectly reasonable for some applications (e.g., three friends choosing a movie list to watch together over the course of a year), but not for others (e.g., when thousands of residents choose a handful of public projects to finance). In fact, it may produce an outcome that may be construed as unfair if it does not reflect the wishes of large groups of voters. Thus, in this work, we address the following question posed by Conitzer et al. [9]:

¹Those familiar with the literature on fair division of private goods will note the conspicuous absence of the *envy freeness* property: that no agent should (strongly) prefer the allocation of another agent. Because we are considering public goods, envy freeness is only vacuously defined: the outcome in our setting is common to all agents.

Is there a stronger fairness notion than proportionality in the public decision making framework...? Although such a notion would not be satisfiable by deterministic mechanisms, it may be satisfied by randomized mechanisms, or it could have novel relaxations that may be of independent interest.

1.3 Summary of Contributions

Our primary contributions are twofold.

- We define a fairness notion for public goods allocation that is stronger than proportionality, ensures fair representation of groups of agents, and in particular, provides a meaningful fairness guarantee even when there are fewer goods than agents.
- We provide polynomial time algorithms for computing integer allocations that approximately satisfy this fairness guarantee for a variety of settings generalizing the public decision making framework and participatory budgeting.

1.4 Core and Approximate Core Outcomes

We define the notion of *core outcomes*, which has been well studied (in similar forms) as a notion of stability in economics [13, 25, 33] and computer science [10, 22] in the context of *randomized* or *fractional* allocations. Our main contribution is to study it in the context of *integer* allocations.

Definition 1.3. Given an outcome \mathbf{c} , we say that a set of agents $S \subseteq N$ form a blocking coalition if there exists an outcome \mathbf{c}' such that $(|S|/n) \cdot u_i(\mathbf{c}') \geq u_i(\mathbf{c})$ for all $i \in S$ and at least one inequality is strict. We say that an outcome \mathbf{c} is a **core outcome** if it admits no blocking coalitions.

Note that non-existence of blocking coalitions of size 1 is equivalent to proportionality, and non-existence of blocking coalitions of size n is equivalent to Pareto optimality. Hence, a core outcome is both proportional and Pareto optimal. However, the core satisfies a stronger property of being, in a sense, Pareto optimal for coalitions of *any size*, provided we scale utilities based on the size of the coalition. Another way of thinking about the core is to view it as a fairness property that enforces a proportionality-like guarantee for coalitions: *e.g.*, if half of all agents have identical preferences, they should be able to get at least half of their maximum possible utility. It is important to note that the core provides a guarantee for every possible coalition. Hence, in satisfying the guarantee for a coalition S , a solution cannot simply make a single member $i \in S$ happy and ignore the rest as this would likely violate the guarantee for the coalition $S \setminus \{i\}$.

Approximate Core. Since a proportional outcome is not guaranteed to exist (even allowing for multiplicative approximations), the same is true for the core. However, an additive approximation to the core still provides a meaningful guarantee, even when there are fewer elements than agents because it provides a non-trivial guarantee to large coalitions of like-minded agents.

Definition 1.4. For $\delta, \alpha \geq 0$, an outcome \mathbf{c} is a **(δ, α) -core outcome** if there exists no set of agents $S \subseteq N$ and outcome \mathbf{c}' such that

$$\frac{|S|}{n} \cdot u_i(\mathbf{c}') \geq (1 + \delta) \cdot u_i(\mathbf{c}) + \alpha$$

for all $i \in S$, and at least one inequality is strict.

A $(0, 0)$ -core outcome is simply a core outcome. A $(\delta, 0)$ -core outcome satisfies δ -proportionality. Similarly, a $(0, 1)$ -core outcome \mathbf{c} satisfies the following relaxation of proportionality that is slightly weaker than proportionality up to one issue: for every agent $i \in N$, $u_i(\mathbf{c}) + 1 \geq Prop_i$. We note that this definition, and by extension, our algorithms satisfy scale invariance, *i.e.*, they are invariant to scaling the utilities of any individual agent. Because we normalize utilities of the agents, the true

additive guarantee is α times the maximum utility an agent can derive from a single element. Since an outcome can have many elements, an approximation with small α remains meaningful.

The advantage of an approximate core outcome is that it fairly reflects the will of a like-minded subpopulation relative to its size. An outcome satisfying approximate proportionality only looks at what *individual* agents prefer, and may or may not respect the collective preferences of subpopulations. We present such an instance in Example 2.1, in effect showing that an approximate core outcome is arguably more fair.

In our results, we will assume $\delta < 1$ to be a small constant, and focus on making α as small as possible. In particular, we desire guarantees on α that exhibit sub-linear or no dependence on n , m , or any other parameters. Deriving such bounds is the main technical focus of our work.

1.5 Our Results

We present algorithms to find approximate core outcomes under matroid, matching, and general packing constraints. Our first result (Section 3) is the following:

THEOREM 1.5. *If feasible outcomes are constrained to be bases of a matroid, then a $(0, 2)$ -core outcome is guaranteed to exist, and for any $\epsilon > 0$, a $(0, 2 + \epsilon)$ -core outcome can be computed in time polynomial in n , m , and $1/\epsilon$.*

In particular, for the public decision making framework, the private goods setting, and multi-winner elections (a.k.a. KNAPSACK with unit sizes), there is an outcome whose guarantee for *every coalition* is close to the guarantee that Conitzer et al. [9] provide to individual agents!

In Section 4, we consider matching constraints. Our result now involves a tradeoff between the multiplicative and additive guarantees.

THEOREM 1.6. *If feasible outcomes are constrained to be matchings in an undirected graph, then for constant $\delta \in (0, 1]$, a $(\delta, 8 + \epsilon/\delta)$ -core outcome can be computed in time polynomial in n and m .*

Our final result (Section 5) is for general packing constraints. Here, our guarantee depends on the width ρ from Equation (1), which captures the difficulty of satisfying the constraints. In particular, the guarantee improves if the constraints are easier to satisfy. This is the most technical result of the paper, and involves different techniques than those used in proving Theorems 1.5 and 1.6; we present an outline of the techniques in Section 5.2.

THEOREM 1.7. *For constant $\delta \in (0, 1)$, given K packing constraints $Ax \leq \mathbf{b}$ with width ρ and $b_k = \omega \left(\frac{\log K}{\delta^2} \right)$ for all $k \in [K]$, there exists a polynomial time computable (δ, α) -core solution, where*

$$\alpha = O \left(\frac{1}{\delta^4} \cdot \log \left(\frac{\min(V_{\max}, n, \rho) \cdot \log^* V_{\max}}{\delta} \right) \right).$$

Here, \log^* is the iterated logarithm, which is the number of times the logarithm function must be iteratively applied before the result becomes less than or equal to 1. Recall that V_{\max} is the maximum utility an agent can have for an outcome (thus $V_{\max} \leq m$); our additive error bound is a vanishing fraction of this quantity. Our bound is also small if the number of agents n is small. Finally, the guarantee improves for small ρ , i.e., as the packing constraints become easier to satisfy. For instance, in participatory budgeting, if the total cost of all projects is only a constant times more than the budget, then our additive guarantee is close to a constant.

Note that V_{\max} (which is bounded by m), n , and ρ are all unrelated quantities. In fact, in Section 5, we state the bound more generally in terms of what we call the *maximally proportionally fair value* R , which informally captures the (existential) difficulty of finding a proportionally fair allocation. The quantity $\min(V_{\max}, n, \rho)$ stems from three different bounds on the value of R .

In the full version [11], we show that the lower bound on \mathbf{b} in the above theorem is necessary: if $\mathbf{b} = O(1)$, then no non-trivial approximation to the core can be guaranteed, even when ρ is a constant. Also in the full version, we consider a different (and more classical) version of the core for general packing constraints, in which a deviating coalition gets a proportional share of resources rather than a proportional share of utility. Our techniques provide a similar approximation to this version of the core, although we do not provide an efficient algorithm in this model.

1.6 Related Work

Core for Public Goods. The notion of core is borrowed from cooperative game theory and was first phrased in game theoretic terms by Scarf [33]. It has been extensively studied in public goods settings [10, 13, 25]. Most literature so far has considered the core with *fractional allocations*. Our definition of core (Definition 1.3) assumes the utility of a deviating coalition is scaled by the size of the coalition. For fractional allocations, one such core allocation coincides with the well-known notion of *proportional fairness*, the extension of the Nash bargaining solution [26]. This solution maximizes the product of the utilities of the agents, and we present the folklore proof in Section 2.2. Our main focus is on finding *integer* allocations that approximate the core, and to the best of our knowledge, this has not been studied previously.

A simpler property than the core is proportionality, which like the core, is impossible to satisfy to any multiplicative approximation using integral allocations. To address this problem, Conitzer et al. [9] defined proportionality up to one issue in the public decision making framework, inspired by related notions for private goods. This guarantee is satisfied by the integral outcome maximizing the *Nash welfare* objective, which is the *geometric* mean of the utilities to the agents. For public goods, this objective is not only NP-HARD to approximate to any multiplicative factor, but approximations to the objective also do not necessarily retain the individual fairness guarantees.

We extend the notion of additive approximate proportionality to additive approximate core outcomes, which provides meaningful guarantees even when there are fewer goods than agents. Unlike proportionality, we show in Section 2.1 that the approach of computing the optimal integral solution to the Nash welfare objective fails to provide a reasonable approximation to the core. Therefore, for our results about matroid constraints (Theorem 1.5) and matching constraints (Theorem 1.6), we slightly modify the integer Nash welfare objective and add a suitable constant term to the utility of each agent. We show that maximizing this smooth objective function achieves a good approximation to the core. However, maximizing this objective is still NP-HARD [12], so we devise local search procedures that run in polynomial time and still give good approximations of the core. In effect, we make a novel connection between appropriate *local optima* of smooth Nash Welfare objectives and the core.

Fairness on Endowments. Classically, the core is defined in terms of endowments, not scaled utilities. In Definition 1.3, we assumed that when a subset S of agents deviates, they can choose any feasible outcome; however, their utility is reduced by a factor that depends on $|S|$. A different notion of core is based on endowments [13, 33] and has been considered in the context of participatory budgeting [10] and in *proportional representation* of voters in multi-winner elections with approval voting. In this notion, a deviating coalition gets a proportional share of resources rather than a proportional share of utility. For example, if the elements have different sizes, and we need to select a subset of them with total size at most B , then a deviating coalition S would get to choose an outcome with size at most $B|S|/n$ instead of B , but would not have its utility scaled down. This notion builds on the seminal work of Foley [13] on the Lindahl equilibrium, from which it can be shown that such a core outcome always exists when fractional allocations are allowed. However, it is not known how to compute such a core outcome efficiently, and it is difficult to

define endowments in settings such as matroid or matching constraints. For integer allocations with packing constraints, we extend our techniques to provide approximations to the notion of core with endowments in the full version [11].

The notion of core with endowments logically implies some fairness notions considered in multi-winner election literature, such as justified representation, extended justified representation [3], and proportional justified representation [32]. Approval-based multi-winner elections are a special case of packing constraints, in which voters (agents) have binary utilities over a pool of candidates (elements), and we must select a set of at most B candidates. The idea behind proportional representation is to define a notion of large cohesive groups of agents with similar preferences, and ensure that such coalitions are proportionally represented. The core on endowments represents a more general condition that holds for all coalitions of agents, not just those that are large and cohesive. Nevertheless, our local search algorithms for Theorems 1.5 and 1.6 are similar to local search algorithms for proportional approval voting (PAV) [4, 35] that achieve proportional representation. It would be interesting to explore the connection between these notions in greater depth.

Private Goods and Envy-freeness. Private goods are a special case of public goods with matroid constraints, and the fair allocation of private goods is well studied [15, 20, 27, 36]. A common fairness criterion for is envy-freeness: that no agent should (strongly) prefer the allocation of another agent. For integral allocations, envy-free allocations or multiplicative approximations thereof may not exist in general. Recent work has introduced envy-freeness up to one good [7, 8, 23, 29], an additive approximation of envy-freeness. The notion of envy does not extend as is to public goods, and the core can be thought of as enforcing envy-freeness across demographics. We note that in addition to resource allocation, group based fairness is also appearing as a desideratum in machine learning. Specifically, related notions may provide a tool against gerrymandered classifiers that appear fair on small samples, but not on structured subsets [19].

Strategyproofness. In this work, we do not consider game-theoretic incentives for manipulation for two reasons. First, even for the restricted case of private goods allocation, preventing manipulation leads to severely restricted mechanisms. For instance, Schummer [34] shows that the only strategyproof and Pareto efficient mechanisms are dictatorial, even when there are only two agents with additive utilities over divisible goods. Second, our work is motivated by public goods settings with many agents, such as participatory budgeting, wherein individual agents often have limited influence over the final outcome as formalized by notions like *strategyproofness in the large* [2].

2 PRELUDE: NASH SOCIAL WELFARE

Our approach to computing approximate core solutions revolves around the Nash social welfare, which is the product (or equivalently, the sum of logarithms) of agent utilities. This objective is commonly considered to be a natural tradeoff between the fairness-blind utilitarian social welfare objective (maximizing the sum of agent utilities) and the efficiency-blind egalitarian social welfare objective (maximizing the minimum agent utility). This function also has the advantage of being *scale invariant* with respect to the utility function of each agent, and in general, preferring more equal distributions of utility.

2.1 Integer Nash Welfare and Smooth Variants

The integer *Max Nash Welfare* (MNW) solution [8, 9] is an outcome \mathbf{c} that maximizes $\sum_{i \in N} \ln u_i(\mathbf{c})$. More technically, if every integer allocation gives zero utility to at least one agent, the MNW solution first chooses a largest set S of agents that can be given non-zero utility simultaneously, and maximizes the product of utilities to agents in S .

Conitzer et al. [9] argue that this allocation is reasonable by showing that it satisfies proportionality up to one issue for public decision making. A natural question is whether it also provides an approximation of the core. *We answer this question in the negative.* The example below shows that even for public decision making (a special case of matroid constraints), the integer MNW solution may fail to return a (δ, α) -core outcome, for any $\delta = o(m)$ and $\alpha = o(m)$.

Example 2.1. Consider an instance of public decision making [9] with m issues and two alternatives per issue. Specifically, each issue t has two alternatives $\{a_1^t, a_2^t\}$, and exactly one of them needs to be chosen. There are two sets of agents $X = \{1, \dots, m\}$ and $Y = \{m+1, \dots, 2m\}$. Every agent $i \in X$ has $u_i^i(a_1^i) = 1$, and utility 0 for all other alternatives. Every agent $i \in Y$ has $u_i^t(a_2^t) = 1$ and $u_i^t(a_1^t) = 1/m$ for all issues $i \in \{1, 2, \dots, m\}$. Visually, this is represented as follows.

	a_1^1	a_2^1	a_1^2	a_2^2	...	a_1^m	a_2^m
$u_{1 \in X}$	1	0	0	0	...	0	0
$u_{2 \in X}$	0	0	1	0	...	0	0
\vdots	\vdots	\vdots	\vdots	\vdots		\vdots	\vdots
$u_{m \in X}$	0	0	0	0	...	1	0
$u_{i \in Y}$	$1/m$	1	$1/m$	1	...	$1/m$	1

The integer MNW outcome is $\mathbf{c} = (a_1^1, a_2^1, \dots, a_1^m)$ because any other outcome gives zero utility to at least one agent. However, coalition Y can deviate, choose outcome $\mathbf{c}' = (a_2^1, a_2^2, \dots, a_2^m)$, and achieve utility m for each agent in Y . For \mathbf{c} to be a (δ, α) -core outcome, we need

$$\exists i \in Y : (1 + \delta) \cdot u_i(\mathbf{c}) + \alpha \geq \frac{|Y|}{|Y| + |X|} \cdot u_i(\mathbf{c}') \quad \Rightarrow \quad 1 + \delta + \alpha \geq \frac{m}{2}.$$

Hence, \mathbf{c} is not a (δ, α) -core outcome for any $\delta = o(m)$ and $\alpha = o(m)$. In contrast, it is not hard to see that \mathbf{c}' is a $(0, 1)$ -core outcome because each agent in X gets utility at most one in any outcome.

Further, note that outcome \mathbf{c} gives every agent utility 1. Since $Prop_i \leq 1$ for each agent i , \mathbf{c} satisfies proportionality, and yet fails to provide a reasonable approximation to the core. One may argue that \mathbf{c}' , which is a $(0, 1)$ -core outcome, is indeed fairer because it respects the utility-maximizing choice of half of the population; the other half of the population cannot agree on what they want, so respecting their top choice is arguably a less fair outcome. Hence, the example also shows that outcomes satisfying proportionality (or proportionality up to one issue) can be very different from and less fair than approximate core outcomes.

Smooth Nash Welfare. One issue with the Nash welfare objective is that it is sensitive to agents receiving zero utility. We therefore consider the following smooth Nash welfare objective:

$$F(\mathbf{c}) := \sum_{i \in N} \ln(\ell + u_i(\mathbf{c})) \quad (2)$$

where $\ell \geq 0$ is a parameter. Note that $\ell = 0$ coincides with the Nash welfare objective. The case of $\ell = 1$ was considered by Fluschnik et al. [12], who showed it is NP-HARD to optimize. Recall that we normalized agent utilities so that each agent has a maximum utility of 1 for any element, so when we add ℓ to the utility of agent i , it is equivalent to adding $\ell \max_j u_{ij}$ to the utility of agent i when utilities are not normalized.

We show that local search procedures for the smooth Nash welfare objective, for appropriate choices of ℓ , yield a $(0, 2)$ -core outcome for matroid constraints (Section 3) and a $(\delta, O(\frac{1}{\delta}))$ -core outcome for matching constraints (Section 4). For packing constraints in contrast, we show in the full version [11] that optimizing any fixed smooth Nash welfare objective cannot guarantee a good approximation to the core, even with a *single* packing constraint, motivating the need for a different algorithm.

2.2 Fractional Max Nash Welfare Solution

For general packing constraints, we use a fractional relaxation of the Nash welfare program. A *fractional outcome* consists of a vector \mathbf{w} such that $w_j \in [0, 1]$ measures the fraction of element j chosen. The utility of agent i under this outcome is $u_i(\mathbf{w}) = \sum_{j=1}^m w_j u_{ij}$. The fractional *Max Nash Welfare* (MNW) solution is a fractional allocation that maximizes the Nash welfare objective (without any smoothing). Define the packing polytope as:

$$\mathcal{P} = \{ \mathbf{w} \in [0, 1]^m \mid \sum_{j=1}^m a_{kj} w_j \leq b_k, \forall k \in [K] \}$$

Then the fractional MNW solution is $\arg \max_{\mathbf{c} \in \mathcal{P}} \sum_i \ln u_i(\mathbf{c})$.

It is easy to show that the fractional MNW allocation lies in the core. Let \mathbf{c} denote the optimal fractional allocation to the MNW program. By first order optimality, for any other allocation \mathbf{d} ,

$$\nabla \ln \mathbf{u}(\mathbf{c}) \cdot (\mathbf{u}(\mathbf{d}) - \mathbf{u}(\mathbf{c})) \leq 0 \quad \Rightarrow \quad \sum_{i \in N} \frac{1}{u_i(\mathbf{c})} (u_i(\mathbf{d}) - u_i(\mathbf{c})) \leq 0 \quad \Rightarrow \quad \sum_{i \in N} \frac{u_i(\mathbf{d})}{u_i(\mathbf{c})} \leq n. \quad (3)$$

Suppose for contradiction that \mathbf{c} is not a core outcome. Then there exists a set of agents $S \subseteq N$ and an outcome \mathbf{d} such that $u_i(\mathbf{d}) \geq (n/|S|) \cdot u_i(\mathbf{c})$, and at least one inequality is tight. This implies $\sum_{i \in S} u_i(\mathbf{d})/u_i(\mathbf{c}) > n$. However, this contradicts Equation (3). Thus \mathbf{c} , the optimal fractional solution to the MNW program, is a core solution.

For the allocation of public goods, it can be shown that the fractional MNW outcome can be irrational despite rational inputs [1], preventing an exact algorithm. For our approximation results, a fractional solution that approximately preserves the utility to each agent would suffice, and we prove the following theorem in the full version [11].

THEOREM 2.2. *For any $\epsilon, \delta > 0$, we can compute a fractional (δ, ϵ) -core outcome in time polynomial in the input size and $\log \frac{1}{\epsilon\delta}$.*

3 MATROID CONSTRAINTS

We now consider public goods allocation with matroid constraints. In particular, we show that when the feasibility constraints encode independent sets of a matroid \mathcal{M} , maximizing the smooth Nash welfare objective in Equation (2) with $\ell = 1$ yields a $(0, 2)$ -core outcome. However, optimizing this objective is known to be NP-HARD [12]. We also show that given $\epsilon > 0$, a local search procedure for this objective function (given below) yields a $(0, 2 + \epsilon)$ -core outcome in polynomial time, which proves Theorem 1.5.

3.1 Algorithm

Fix $\epsilon > 0$. Let $\gamma = \frac{\epsilon}{4m}$, where $m = |W|$ is the number of elements. Recall that there are n agents.

- (1) Start with an arbitrary basis \mathbf{c} of \mathcal{M} .
- (2) Compute $F(\mathbf{c}) = \sum_{i \in N} \ln(1 + u_i(\mathbf{c}))$.
- (3) Let a *swap* be a pair (j, j') such that $j \in \mathbf{c}$, $j' \notin \mathbf{c}$, and $\mathbf{c}' = \mathbf{c} \setminus \{j\} \cup \{j'\}$ is also a basis of \mathcal{M} .
- (4) Find a swap such that $F(\mathbf{c}') - F(\mathbf{c}) \geq \frac{\gamma}{m}$.
 - If such a swap exists, then perform the swap, i.e., update $\mathbf{c} \leftarrow \mathbf{c}'$, and go to Step (2).
 - If no such swap exists, then output \mathbf{c} as the final outcome.

3.2 Analysis

First, we show that the local search algorithm runs in time polynomial in n , m , and $1/\epsilon$. Note that $F(\mathbf{c}) = O(n \ln m)$ because in our normalization, each agent can have utility at most m . Thus, the number of iterations is $O(m^2 \log m/\epsilon)$. Finally, each iteration can be implemented in $O(n \cdot m^2)$ time by iterating over all pairs and computing the change in the smooth Nash welfare objective.

Next, let \mathbf{c}^* denote the outcome maximizing the smooth Nash welfare objective with $\ell = 1$, and $\widehat{\mathbf{c}}$ denote the outcome returned by the local search algorithm. We show that \mathbf{c}^* is a $(0, 2)$ -core outcome, while $\widehat{\mathbf{c}}$ is a $(0, 2 + \epsilon)$ -core outcome.

For outcome \mathbf{c} , define $F_i(\mathbf{c}) = \ln(1 + u_i(\mathbf{c}))$. Fix an arbitrary outcome \mathbf{c} . For an agent i with $u_i(\mathbf{c}) > 0$, we have that for every element $j \in \mathbf{c}$:

$$F_i(\mathbf{c}) - F_i(\mathbf{c} \setminus \{j\}) \leq \frac{u_{ij}}{u_i(\mathbf{c}) + 1 - u_{ij}} \leq \frac{u_{ij}}{u_i(\mathbf{c})}.$$

This holds because $\ln(x + h) - \ln x \leq \frac{h}{x}$ for $x > 0$ and $h \geq 0$. Summing this over all $j \in \mathbf{c}$ gives

$$\sum_{j \in \mathbf{c}} F_i(\mathbf{c}) - F_i(\mathbf{c} \setminus \{j\}) \leq \sum_{j \in \mathbf{c}} \frac{u_{ij}}{u_i(\mathbf{c})} = \frac{u_i(\mathbf{c})}{u_i(\mathbf{c})} = 1.$$

For an agent i with $u_i(\mathbf{c}) = 0$, we trivially have $\sum_{j \in \mathbf{c}} F_i(\mathbf{c}) - F_i(\mathbf{c} \setminus \{j\}) = 0$. Summing over all agents, we have that for every outcome \mathbf{c} :

$$\sum_{j \in \mathbf{c}} F(\mathbf{c}) - F(\mathbf{c} \setminus \{j\}) = \sum_{i \in N} \sum_{j \in \mathbf{c}} F_i(\mathbf{c}) - F_i(\mathbf{c} \setminus \{j\}) \leq n. \quad (4)$$

We now use the following result:

LEMMA 3.1 ([21]). *For every pair of bases \mathbf{c} and \mathbf{c}' of a matroid \mathcal{M} , there is a bijection $f : \mathbf{c} \rightarrow \mathbf{c}'$ such that for every $j \in \mathbf{c}$, $\mathbf{c} \setminus \{j\} \cup \{f(j)\}$ is also a basis.*

Using the above lemma, combined with the fact that $\ln(x + h) - \ln x \geq \frac{h}{x+h}$ for $x > 0$ and $h \geq 0$, we have that for all \mathbf{c}, \mathbf{c}' :

$$\begin{aligned} \sum_{j \in \mathbf{c}} F(\mathbf{c} \setminus \{j\} \cup \{f(j)\}) - F(\mathbf{c} \setminus \{j\}) &\geq \sum_i \sum_{j \in \mathbf{c}} \frac{u_{if(j)}}{u_i(\mathbf{c}) + 1 - u_{ij} + u_{if(j)}} \\ &\geq \sum_{i \in S} \sum_{j' \in \mathbf{c}'} \frac{u_{ij'}}{u_i(\mathbf{c}) + 2} = \sum_{i \in S} \frac{u_i(\mathbf{c}')}{u_i(\mathbf{c}) + 2}. \end{aligned} \quad (5)$$

We now provide almost similar proofs for the approximations achieved by the global optimum \mathbf{c}^* and the local optimum $\widehat{\mathbf{c}}$.

Global optimum. Suppose for contradiction that \mathbf{c}^* is not a $(0, 2)$ -core outcome. Then, there exist a subset S of agents and an outcome \mathbf{c}' such that for all $i \in S$,

$$\frac{|S|}{n} \cdot u_i(\mathbf{c}') \geq u_i(\mathbf{c}^*) + 2,$$

and at least one inequality is strict. Rearranging the terms and summing over all $i \in S$, we obtain:

$$\sum_{i \in S} \frac{u_i(\mathbf{c}')}{u_i(\mathbf{c}^*) + 2} > \sum_{i \in S} \frac{n}{|S|} = n.$$

Combining this with Equation (5), and subtracting Equation (4) yields:

$$\sum_{j \in \mathbf{c}^*} (F(\mathbf{c}^* \setminus \{j\} \cup \{f(j)\}) - F(\mathbf{c}^*)) > 0.$$

This implies existence of a pair $(j, f(j))$ such that $F(\mathbf{c}^* \setminus \{j\} \cup \{f(j)\}) - F(\mathbf{c}^*) > 0$, which contradicts the optimality of \mathbf{c}^* because $\mathbf{c}^* \setminus \{j\} \cup \{f(j)\}$ is also a basis of \mathcal{M} .

Local optimum. Similarly, suppose for contradiction that $\widehat{\mathbf{c}}$ is not a $(0, 2 + \epsilon)$ -core outcome. Then, there exist a subset S of agents and an outcome \mathbf{c}' such that for all $i \in S$,

$$\frac{|S|}{n} \cdot u_i(\mathbf{c}') \geq u_i(\widehat{\mathbf{c}}) + 2 + \epsilon > (1 + \gamma)(u_i(\widehat{\mathbf{c}}) + 2).$$

Here, the final transition holds because $\gamma < \epsilon/(m+2) \leq \epsilon/(u_i(\widehat{\mathbf{c}}) + 2)$. Again, rearranging and summing over all $i \in S$, we obtain:

$$\sum_{i \in S} \frac{u_i(\mathbf{c}')}{u_i(\widehat{\mathbf{c}}) + 2} > (1 + \gamma) \sum_{i \in S} \frac{n}{|S|} \geq n \cdot (1 + \gamma).$$

Once again, combining this with Equation (5), and subtracting Equation (4) yields:

$$\sum_{j \in \widehat{\mathbf{c}}} (F(\widehat{\mathbf{c}} \setminus \{j\} \cup \{f(j)\}) - F(\widehat{\mathbf{c}})) > n\gamma.$$

This implies existence of a pair $(j, f(j))$ such that $F(\widehat{\mathbf{c}} \setminus \{j\} \cup \{f(j)\}) - F(\widehat{\mathbf{c}}) > n\gamma/m$, which violates local optimality of $\widehat{\mathbf{c}}$ because $\widehat{\mathbf{c}} \setminus \{j\} \cup \{f(j)\}$ is also a basis of \mathcal{M} .

Lower Bound. While a $(0, 2)$ -core always outcome exists, we show in the following example that a $(0, 1 - \epsilon)$ -core outcome is not guaranteed to exist for any $\epsilon > 0$.

LEMMA 3.2. *For $\epsilon > 0$ and matroid constraints, $(0, 1 - \epsilon)$ -core outcomes are not guaranteed to exist.*

PROOF. Consider the following instance of public decision making where we have several issues and must choose a single alternative for each issue, a special case of matroid constraints. There are n agents, where n is even. There are $m = (n - 2) + n/2$ issues. The first $n - 2$ issues correspond to unit-value private goods, *i.e.*, each such issue has n alternatives, and each alternative gives utility 1 to a unique agent and utility 0 to others. The remaining $n/2$ issues are “pair issues”; each such issue has $\binom{n}{2}$ alternatives, one corresponding to every pair of agents that gives both agents in the pair utility 1 and all other agents utility 0.

It is easy to see that every integer allocation gives utility at most 1 to at least two agents. Consider the deviating coalition consisting of these two agents. They can choose the alternative that gives them each utility 1 on every pair issue, and split the $n - 2$ private good equally. Thus, they each get utility $n/2 + (n - 2)/2 = n - 1$. For the outcome to be a $(0, \alpha)$ -core outcome, we need $1 + \alpha \geq (2/n) \cdot (n - 1)$. As $n \rightarrow \infty$, this requires $\alpha \rightarrow 1$. Hence, for any $\epsilon > 0$, a $(0, 1 - \epsilon)$ -core outcome is not guaranteed to exist. \square

Note that Theorem 1.5 shows existence of a $(0, 2)$ -core outcome, which is therefore tight up to a unit additive relaxation. Whether a $(0, 1)$ -core outcome always exists under matroid constraints remains an important open question. Interestingly, we show that such an outcome always exists for the special case of private goods allocation, and, in fact, can be achieved by maximizing the smooth Nash welfare objective. We provide a proof of this result in the full version [11].

LEMMA 3.3. *For private goods allocation, maximizing the smooth Nash welfare objective with $\ell = 1$ returns a $(0, 1)$ -core outcome.*

4 MATCHING CONSTRAINTS

We now present the algorithm proving Theorem 1.6. We show that if the elements are edges of an undirected graph $G(V, E)$, and the feasibility constraints encode a matching, then for constant $\delta \in (0, 1]$, a $(\delta, 8 + \frac{6}{\delta})$ -core always exists and is efficiently computable. The idea is to again run a *local search* on the smooth Nash welfare objective in Equation (2), but this time with $\ell \approx 1 + \frac{4}{\delta}$.

Algorithm. Recall that there are n agents. Let $|V| = r$ and $|E| = m$. Let $\kappa = \frac{2}{\delta}$. For simplicity, assume $\kappa \in \mathbb{N}$. Our algorithm is inspired by the PRAM algorithm for approximate maximum weight matchings due to Hougardy and Vinkemeier [18], and we follow their terminology. Given a matching \mathbf{c} , an *augmentation* with respect to \mathbf{c} is a matching $T \subseteq E \setminus \mathbf{c}$. The *size* of the augmentation

is $|T|$. Let $M(T)$ denote the subset of edges of \mathbf{c} that have a vertex which is matched under T . Then, the matching $(\mathbf{c} \setminus M(T)) \cup T$ is called the *augmentation* of \mathbf{c} using T .

- (1) Start with an arbitrary matching \mathbf{c} of G .
- (2) Compute $F(\mathbf{c}) = \sum_i \ln(1 + 2\kappa + u_i(\mathbf{c}))$.
- (3) Let \mathcal{C} be the set of all augmentations with respect to \mathbf{c} of size at most κ .
 - If there exists $T \in \mathcal{C}$ such that $F((\mathbf{c} \setminus M(T)) \cup T) - F(\mathbf{c}) \geq \frac{\eta}{\kappa r}$, perform this augmentation (i.e., let $\mathbf{c} \leftarrow (\mathbf{c} \setminus M(T)) \cup T$) and go to Step (2).
 - Otherwise, output \mathbf{c} as the final outcome.

Analysis and Lower Bound. Since the outline of the analysis is very similar to the analysis for matroid constraints, we provide the proof of Theorem 1.6 in the full version [11]. Also in the full version, we show that for matching constraints, a $(\delta, 1 - \epsilon)$ -core is not guaranteed to exist for any $\delta \geq 0$ and $\epsilon > 0$. This leaves the interesting open question of whether a $(0, 1)$ -core solution always exists for matching constraints.

5 GENERAL PACKING CONSTRAINTS

In this section, we study approximation to the core under general packing constraints of the form $A\mathbf{x} \leq \mathbf{b}$. Recall that there are m elements, V_i is the maximum possible utility that agent i can receive from a feasible outcome, and $V_{\max} = \max_{i \in N} V_i$. We prove a statement slightly more general than Theorem 1.7. We first need the following concept.

5.1 Maximal Proportionally Fair Outcome

Given an instance of public goods allocation subject to packing constraints, we define the notions of an r -proportionally fair (r -PF) outcome, a maximally proportionally fair (MPF) outcome, and the MPF value of the instance.

Definition 5.1 (MPF Outcome). For $r > 0$, we say that a fractional outcome \mathbf{w} is r -proportionally fair (r -PF) if it satisfies:

$$u_i(\mathbf{w}) \geq \frac{V_i}{r} - 1, \quad \forall i \in N.$$

The *maximally proportionally fair* (MPF) value R of an instance is the least value r such that there exists an r -PF outcome. For simplicity, we say that an R -PF outcome is a *maximally proportionally fair* (MPF) outcome.

This concept is crucial to stating and deriving our approximation results. An r -PF outcome gives each agent an r fraction of its maximum possible utility V_i (which can be thought of as the fair share guarantee of the agent), if the agent is given 1 unit of utility for free. Thus, a smaller value of r indicates a better solution. The MPF value R denotes the best possible guarantee. The additive 1 in Def. 5.1 can be replaced by any positive constant; we choose 1 for simplicity. We give an upper bound for R that holds for all instances. Recall from Equation (1) that ρ is the *width* of the instance.

LEMMA 5.2. $R \leq \min(V_{\max}, n, \rho)$, and an MPF outcome is computable in polynomial time.

PROOF. To show that R is well-defined, note that for $r = V_{\max}$, an r -PF outcome \mathbf{w} simply requires $u_i(\mathbf{w}) \geq 0$, which is trivially achieved by every outcome. Therefore, R is well-defined, and $R \leq V_{\max}$. Next, $R \leq n$ follows from the fact that there exist fractional outcomes satisfying proportionality (e.g., the outcome \mathbf{w} obtained by taking the uniform convex combination of the n outcomes that give optimal to each individual agent). Finally, to show $R \leq \rho$, consider the outcome \mathbf{w} in which $w_j = \frac{1}{\rho}$ for each element j . Clearly, $u_i(\mathbf{w}) \geq \frac{V_i}{\rho}$ for all i . Further, $A\mathbf{w} \leq \mathbf{b}$ is satisfied trivially due to the fact that ρ is the width of the packing constraints.

To compute the value of R as well as an MPF outcome, we first note that the value of V_i for each agent i can be computed by solving a separate LP. Then, we consider the following LP:

$$\begin{aligned} & \text{Maximize } \hat{r} && (6) \\ & \sum_{j \in W} u_{ij} w_j \geq V_i \cdot \hat{r} - 1 \quad \forall i \in [n] \\ & \mathbf{Aw} \leq \mathbf{b} \\ & w_j \in [0, 1] \quad \forall j \in W \end{aligned}$$

Here, $\mathbf{Aw} \leq \mathbf{b}$ are the packing constraints of the instance, and \hat{r} is a variable representing $1/r$. Thus, maximizing \hat{r} minimizes r , which yields an MPF outcome. This can be accomplished by solving $n + 1$ linear programs, which can be done in polynomial time. \square

Our main result in this section uses any r -PF outcome, and provides a guarantee in terms of $\log r$. Thus, we do not need to necessarily compute an exact MPF outcome. We note that an MPF outcome can be very different from a core outcome. Yet, an MPF outcome gives each agent a large fraction of its maximum possible utility, subject to a small additive relaxation. As we show below, this helps us find integral outcomes that provide good approximations of the core.

5.2 Result and Proof Idea

Our main result for this section (Theorem 1.7) can be stated in a refined way as follows. Recall that \log^* is the iterated logarithm, which is the number of times the logarithm function must be iteratively applied before the result becomes less than or equal to 1.

THEOREM 5.3. *Fix constant $\delta \in (0, 1)$. Suppose we are given a set of K packing constraints $\mathbf{Ax} \leq \mathbf{b}$ such that $b_k = \omega\left(\frac{\log K}{\delta^2}\right)$ for all $k \in [K]$. Let R be the MPF value of this instance. Then there exists a polynomial time computable (δ, α) -core outcome, where*

$$\alpha = O\left(\frac{1}{\delta^4} \cdot \log\left(\frac{R \cdot \log^* V_{\max}}{\delta}\right)\right).$$

We first note that the above result cannot be obtained by maximizing the smooth Nash welfare objective as we prove in the full version of this paper [11]. To be precise, the example shows that no single value of parameter ℓ in the smooth Nash welfare objective can provide a polylog additive guarantee for all instances. While it may be possible to choose the value of ℓ based on the instance, it does not seem trivial. We take a different approach. Our idea is to start with a fractional core solution \mathbf{x} . Suppose it assigns utility U_i^* to agent i . Fix $\delta > 0$, and consider the following program.

$$\begin{aligned} & \text{Minimize } \alpha && (7) \\ & \alpha + (1 + \delta) \cdot \sum_{j \in W} u_{ij} w_j \geq U_i^* \quad \forall i \in [n] \\ & \mathbf{Aw} \leq \mathbf{b} \\ & w_j \in \{0, 1\} \quad \forall j \in W \\ & \alpha \geq 0 \end{aligned}$$

For the optimum value α^* , we obtain an outcome that is (δ, α') -core for every $\alpha' > \alpha$. To see this, take a subset of agents S and a feasible utility vector \mathbf{U}' under any other (even fractional) outcome. Because \mathbf{x} is a core outcome, there exists $i \in S$ such that $U_i^* \geq (|S|/n) \cdot U_i'$. For $\alpha' > \alpha$, the ILP solution implies

$$\alpha' + (1 + \delta) \cdot \sum_{j \in W} u_{ij} w_j > U_i^* \geq \frac{|S|}{n} \cdot U_i',$$

which implies that the solution is (δ, α') -core according to Definition 1.4.

However, α^* obtained from this program can be rather large, as illustrated in the following example. Consider the **KNAPSACK** setting with m unit-size projects. There is an overall budget $B = m/2$. For every feasible integral outcome \mathbf{c} , let there be an agent with utility 1 for every project in \mathbf{c} and 0 for all other projects. Thus, there are $\binom{m}{m/2}$ agents. The fractional core outcome gives weight $1/2$ to each project, thus giving utility $V_i/2 = m/4$ to each agent i . However, every integral outcome gives utility 0 to at least one agent, which implies $\alpha^* = \Omega(m)$.

This example shows that when there are a large number of agents, we cannot achieve Theorem 5.3 by hoping to approximately preserve the utilities to *all* agents with respect to the fractional core solution. However, note that in the above example, though there is one agent who gets very little utility, this agent has no incentive to deviate if she is given one unit of utility for free. This insight leads us to our analysis below, which is based on rounding the fractional core solution \mathbf{x} .

Let us apply randomized rounding to \mathbf{x} . Instead of using Chernoff bounds to ensure that there are no “violations” (i.e., that no agent receives utility that is too far from its utility under the core outcome \mathbf{x}), we hope to bound the expected number of such violations. If there are few such agents, we still have an approximate core outcome because if this small coalition of agents deviates, its utility under a new outcome will be scaled down by a large factor. Unfortunately, it can be shown that bounding the expected number of deviations by a sufficiently small number forces $\alpha = \Omega(\log V_{\max})$. This is better than $\alpha = \Omega(m)$ from our previous approach, but still *much* larger than the bound we want to achieve in Theorem 5.3 when the width ρ is small.

This brings up the main technical idea. We observe that an MPF outcome, though not in the core, provides a reasonably large utility to each agent. We add a small amount of this outcome to the fractional core before applying randomized rounding. We are now ready to present our algorithm.

5.3 Algorithm

Fix $\delta \in (0, 1)$, and let $\gamma = \frac{\delta}{8}$.

- (1) Compute the (approximate) fractional core solution \mathbf{x} as in Theorem 2.2, where x_j is the fraction of element j chosen.
- (2) Let \mathbf{y} be an MPF outcome as in Definition 5.1.
- (3) Let $\mathbf{z} = (1 - \gamma)\mathbf{x} + \gamma\mathbf{y}$.
- (4) For each $j \in W$, choose j to be in the outcome \mathbf{c} independently with probability $\hat{z}_j = (1 - \gamma)z_j$.

5.4 Analysis

We show that this algorithm yields, with at least a constant probability, a feasible outcome that satisfies the guarantee in Theorem 5.3. This directly shows the existence of such an outcome. Note that the fractional Max Nash Welfare solution \mathbf{x} can be irrational, but we can compute an approximation in polynomial time (see Theorem 2.2 for details), which does not change our guarantee asymptotically. Further, \mathbf{y} can be computed in polynomial time (Lemma 5.2). Hence, the algorithm runs in expected polynomial time.

We first show that the packing constraints are satisfied. Since we scale down \mathbf{z} by a factor $(1 - \gamma)$ before rounding, we have $A\hat{\mathbf{z}} \leq (1 - \gamma)\mathbf{b}$. Since $\mathbf{b} = \omega\left(\frac{\log K}{\delta^2}\right)$, a simple application of Chernoff bounds shows that with probability at least 0.99, the rounded solution \mathbf{c} satisfies $A\mathbf{c} \leq \mathbf{b}$. Therefore, if we show that the algorithm also yields the desired approximation of the core with at least a constant probability (1/6 to be precise), we will be done by applying union bound to the two events, feasibility and good approximation to the core.

For the ease of presentation, we suppress constants throughout the proof and use the asymptotic notation liberally. We also assume that $V_{\max} = \omega(1)$ since otherwise there is a trivial $(0, O(1))$ -core

outcome that chooses a null outcome, giving zero utility to each agent. For lack of space, we relegate parts of the proof to the full version of this paper [11].

5.4.1 Grouping Agents. In order to analyze our algorithm, we partition the agents into groups with exponentially decreasing values of V_i . Recall that V_i is the maximum utility that agent i can get from any outcome. Set $Q_0 = \log V_{\max}$, and for $\ell = 0, 1, \dots, L-1$, define group G_ℓ as:

$$G_\ell = \{i \in N \mid Q_\ell \geq \log V_i \geq Q_{\ell+1}\}.$$

Here, for $\ell = 0, 1, \dots, L-1$, we define: $Q_{\ell+1} = 2 \log Q_\ell$.

We call G_0, \dots, G_{L-1} the *heavy groups*. We choose L so that $Q_L = \Theta\left(\log \frac{R \log^* V_{\max}}{\gamma^3}\right)$. This implies $L = \Omega(\log^* V_{\max}) = \omega(1)$, since $V_{\max} = \omega(1)$. For agent i in a heavy group, $V_i \geq e^{Q_\ell} \geq \frac{2RL}{\gamma^3} > 2R$. Thus, the utility that the MPF solution provides to agent i is at least $\frac{V_i}{R} - 1 \geq \frac{V_i}{2R}$.

Finally, we put the remaining agents (with a small V_i) in a *light group* defined as follows:

$$G_L = \{i \in N \mid \log V_i \leq Q_L\}.$$

The MPF solution may not provide any guarantee for the utility of agents in this group.

5.4.2 Bounding Violations of Utility Preservation. We want to bound the number of agents whose utilities are far from those under the core outcome. First, we need a specialized Chernoff bound.

LEMMA 5.4. (Proved in full version) *Let X_1, X_2, \dots, X_q be independent random variables in $[0, 1]$, and let $X = \sum_{j=1}^q X_j$. For $\gamma \in (0, 1/2)$, suppose $\mathbb{E}[X] = (1 - \gamma) \cdot A + \gamma \cdot B$ for $A, B \geq 0$. Then*

$$\Pr[X < (1 - 2\gamma) \cdot A] \leq e^{-\frac{\gamma^3}{2} \max(B, A/2)}$$

Recall that \mathbf{x} is the fractional MNW solution, \mathbf{y} is the fractional MPF solution, and our algorithm applies randomized rounding to their scaled down mixture $(1 - \gamma)\mathbf{z} = (1 - \gamma)^2\mathbf{x} + \gamma(1 - \gamma)\mathbf{y}$. Let \widehat{U}_i denote the utility of agent i under the final integral outcome obtained by randomly rounding $(1 - \gamma)\mathbf{z}$. Recall that U_i^* is the utility of agent i under the core outcome \mathbf{x} . We want to show that \widehat{U}_i is either multiplicatively or additively close to U_i^* for most agents. For a heavy group G_ℓ , where $\ell \in \{0, 1, \dots, L-1\}$, define

$$F_\ell = \left\{i \in G_\ell \mid \widehat{U}_i < (1 - 3\gamma)U_i^*\right\}.$$

Similarly, for the light group G_L , define

$$F_L = \left\{i \in G_L \mid \widehat{U}_i < \min\left((1 - 3\gamma)U_i^*, U_i^* - \frac{4Q_L}{\gamma^4}\right)\right\}.$$

We will use Lemma 5.4 to bound the sizes of F_ℓ for $\ell \in \{0, 1, \dots, L\}$ as follows.

THEOREM 5.5. (Proved in full version) *We have that:*

- (1) *With probability at least $2/3$, we have $|F_\ell| \leq \frac{1}{2Le^{Q_\ell}} \cdot |G_\ell|$, $\forall \ell \in \{0, 1, \dots, L-1\}$.*
- (2) *With probability at least $1/2$, we have $|F_L| \leq \frac{1}{2e^{Q_L}} \cdot |G_L|$.*

Thus, with probability at least $1/6$, both the above inequalities hold simultaneously.

5.4.3 Approximate Core. We showed that with probability at least $1/6$, our algorithm returns a solution that satisfies conditions in both parts of Theorem 5.5. We now show that such a solution is the desired approximate core solution. The main idea is that when a set of agents deviate, the fraction of agents in a group G_ℓ that are in F_ℓ is small enough such that even if they receive their maximum possible utility, which is e^{Q_ℓ} , their scaled down utility is at most a constant.

THEOREM 5.6. For every coalition S and every possible outcome \mathbf{h} , there exists an agent $i \in S$ s.t.

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq (1 + 8\gamma) \cdot \widehat{U}_i + \frac{5Q_L}{\gamma^4}.$$

PROOF. Let $W = N \setminus \cup_{\ell=0}^L F_\ell$. In other words, W is the set of agents who either receive a good multiplicative approximation to their expected utility in the core (for the heavy groups), or a good additive approximation to their expected utility in the core (for the light group). In particular, for every $i \in W$, we have $\widehat{U}_i \geq \min\left((1 - 3\gamma) \cdot U_i^*, U_i^* - \frac{4Q_L}{\gamma^4}\right)$, which implies

$$U_i^* \leq \frac{1}{1 - 3\gamma} \cdot \widehat{U}_i + \frac{4Q_L}{\gamma^4}. \quad (8)$$

Consider a set of agents S that may want to deviate, and let \mathbf{h} be any (even fractional) outcome. There are two cases:

Case 1. Suppose $|S \cap W| \geq (1 - \gamma) \cdot |S|$. Then, due to the fractional core optimality condition (see Section 2.2), we have:

$$\sum_{i \in S \cap W} \frac{U_i(\mathbf{h})}{U_i^*} \leq n.$$

Note that in polynomial time, Theorem 2.2 only finds an approximate solution whose utilities $\{\tilde{U}_i\}$ with $\sum_{i \in S \cap W} \frac{U_i(\mathbf{h})}{\tilde{U}_i} \leq n(1 + \eta)$ for small $\eta > 0$. It is easy to check this does not alter the rest of the proof and adds a small multiplicative factor of $(1 + \eta)$ to the final approximation bound. We ignore this factor for simplicity and simply assume $\{U_i^*\}$ are the optimal MNW utilities. The above implies $\frac{|S|}{n} \cdot \sum_{i \in S \cap W} \frac{u_i(\mathbf{h})}{U_i^*} \leq |S| \leq \frac{1}{1 - \gamma} \cdot |S \cap W|$.

Therefore, there exists an agent $i \in S \cap W$ such that

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq \frac{1}{1 - \gamma} \cdot U_i^* \leq \frac{1}{1 - \gamma} \cdot \left(\frac{1}{1 - 3\gamma} \cdot \widehat{U}_i + \frac{4Q_L}{\gamma^4} \right),$$

where the last transition is due to Equation (8) and the fact that $i \in W$. Finally, it is easy to check that for $\gamma = \delta/8 \leq 1/8$, we have $\frac{1}{(1-\gamma)(1-3\gamma)} \leq 1 + 8\gamma$ and $4/(1 - \gamma) \leq 5$, which yields:

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq (1 + 8\gamma) \cdot \widehat{U}_i + \frac{5Q_L}{\gamma^4}. \quad (9)$$

Case 2. Otherwise, $|S \setminus W| \geq \gamma|S|$. In this case, we want to show that there exists an agent $i \in S \setminus W$ such that $(|S|/n) \cdot u_i(\mathbf{h}) \leq 1/\gamma$. Because $\widehat{U}_i \geq 0$ and $Q_L = \omega(1)$, such an agent will also satisfy Equation (9). We show this by taking two sub-cases.

First, suppose the light group satisfies $|S \cap F_L| \geq \frac{\gamma}{2}|S|$. Then: $|S| \leq \frac{2}{\gamma} \cdot |S \cap F_L| \leq \frac{2}{\gamma} \cdot |F_L|$. Thus, for any agent $i \in F_L$, we have

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq \frac{2}{\gamma n} \cdot |F_L| \cdot V_i \leq \frac{2}{\gamma n} \cdot \frac{|G_L|}{2e^{Q_L}} \cdot V_i \leq \frac{1}{\gamma}.$$

Here, the second transition follows from Theorem 5.5. To see why the third transition holds, note that $|G_L| \leq n$, and that $\log V_i \leq Q_L$ because $i \in G_L$.

Similarly, in the other sub-case, suppose $|S \cap F_L| \leq \frac{\gamma}{2}|S|$. Then, there exists a heavy group $\ell \in \{0, 1, \dots, L - 1\}$ such that $|S \cap F_\ell| \geq \frac{\gamma}{2L}|S|$. This means $|S| \leq \frac{2L}{\gamma} \cdot |S \cap F_\ell| \leq \frac{2L}{\gamma} \cdot |F_\ell|$.

Again, for an arbitrary agent $i \in F_\ell$, we have:

$$\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq \frac{2L}{\gamma n} \cdot |F_\ell| \cdot V_i \leq \frac{2L}{\gamma n} \cdot \frac{|G_\ell|}{2Le^{Q_\ell}} \cdot V_i \leq \frac{1}{\gamma}.$$

Once again, the third transition follows from Theorem 5.5, and the fourth transition holds because $|G_\ell| \leq n$ and $\log V_i \leq Q_\ell$ as $i \in G_\ell$. Putting everything together, the theorem follows. \square

Since $\gamma = \frac{\delta}{8}$ and $Q_L = \Theta\left(\log \frac{R \log^* V_{\max}}{\gamma}\right)$, Theorem 5.6 implies $\frac{|S|}{n} \cdot u_i(\mathbf{h}) \leq (1 + \delta) \cdot \widehat{U}_i + \alpha^*$, where $\alpha^* = O\left(\frac{1}{\delta^4} \cdot \log\left(\frac{R \cdot \log^* V_{\max}}{\delta}\right)\right)$. The existence of such an agent implies that a solution satisfying Theorem 5.5 is a (δ, α) -core solution for every $\alpha > \alpha^*$, which completes the proof of Theorem 5.3.

6 CONCLUSION

We considered the problem of fairly allocating public goods. We argued that the *core*, which is a generalization of proportionality and Pareto efficiency, and approximations of the core provide reasonable fairness guarantees in this context. Given that no integral outcome may be in the core, we presented efficient algorithms to produce integral outcomes that are constant or near-constant approximations of the core, thereby also establishing the non-trivial existence of such outcomes. Note Our algorithms for matroid and matching constraints that globally optimize the smooth Nash welfare objective achieve *exact* Pareto efficiency, in addition to an approximation of the core. An interesting question is whether the same guarantee can be provided for general packing constraints.

Another natural question following our work is to tighten our upper bounds, or to establish matching lower bounds. We show the existence of a $(0, 2)$ -core outcome for matroid constraints (Theorem 1.5), but our lower bound only shows that a $(0, 1 - \epsilon)$ -core outcome may not exist. This leaves open the question of whether a $(0, 1)$ -core outcome always exists. Existence of $(0, 1)$ -core outcome is also an open question for matching constraints. For packing constraints, it is unknown if even a (δ, α) -core outcome exists for constant $\delta > 0$ and $\alpha = O(1)$. This also remains an open question for the endowment-based notion of core we consider in the full version of this paper [11].

At a higher level, we established connections between approximating the core in our multi-agent environment and the problem of finding the optimal (i.e., utility-maximizing) outcome for a *single* agent. For instance, given matching constraints, our algorithm uses the idea of *short* augmenting paths from fast PRAM algorithms. This hints at the possibility of a deeper connection between efficiency results and existence results.

ACKNOWLEDGMENTS

We thank Vincent Conitzer, Rupert Freeman, Ashish Goel, and Paul Gözl for helpful comments. Brandon Fain is supported by NSF grants CCF-1637397 and IIS-1447554. Kamesh Munagala is supported by NSF grants CCF-1408784, CCF-1637397, and IIS-1447554. Nisarg Shah is supported by the Discovery Grants program from the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, and J. Lang. 2017. Positional Social Decision Schemes: Fair and Efficient Randomized Voting. Manuscript.
- [2] E. M. Azevedo and E. Budish. 2017. *Strategy-proofness in the Large*. Technical Report. National Bureau of Economic Research.
- [3] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [4] H. Aziz, E. Elkind, S. Huang, M. Lackner, L. Sanchez Fernandez, and P. Skowron. 2018. On the Complexity of Extended and Proportional Justified Representation. In *AAAI Conference on Artificial Intelligence*.
- [5] H. Aziz and S. Mackenzie. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Any Number of Agents. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. 416–427.
- [6] G. Benade, S. Nath, A. D. Procaccia, and N. Shah. 2017. Preference Elicitation for Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 376–382.

- [7] E. Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [8] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 305–322.
- [9] V. Conitzer, R. Freeman, and N. Shah. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 629–646.
- [10] B. Fain, A. Goel, and K. Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*. 384–399.
- [11] B. Fain, K. Munagala, and N. Shah. 2018. Fair Allocation of Indivisible Public Goods. *ArXiv e-prints* (May 2018). arXiv:1805.03164
- [12] T. Fluschnik, P. Skowron, M. Triphaus, and K. Wilker. 2017. Fair Knapsack. arXiv:1711.04520.
- [13] D. K. Foley. 1970. Lindahl’s Solution and the Core of an Economy with Public Goods. *Econometrica* 38, 1 (1970), 66–72.
- [14] N. Garg, V. Kamble, A. Goel, D. Marn, and K. Munagala. 2017. Collaborative Optimization for Collective Decision-making in Continuous Spaces. In *Proceedings of the 26th International World Wide Web Conference (WWW)*. 617–626.
- [15] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation (NSDI)*. 24–37.
- [16] A. Goel, A. K. Krishnaswamy, and S. Sakshuwong. 2016. Budget Aggregation via Knapsack Voting: Welfare-maximization and Strategy-proofness. *Collective Intelligence* (2016).
- [17] A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2015. Knapsack voting. *Collective Intelligence* (2015).
- [18] S. Hougardy and D. E. Vinkemeier. 2006. Approximating Weighted Matchings in Parallel. *Inform. Process. Lett.* 99, 3 (2006), 119 – 123.
- [19] M. Kearns, S. Neel, A. Roth, and Z. S. Wu. 2017. Preventing Fairness Gerrymandering: Auditing and Learning for Subgroup Fairness. *ArXiv e-prints* (Nov. 2017). arXiv:cs.LG/1711.05144
- [20] F. P. Kelly, A. K. Maulloo, and D. KH Tan. 1998. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society* (1998), 237–252.
- [21] J. P. S. Kung. 1986. Basis-Exchange Properties. In *Theory of Matroids*, White N. (Ed.). Cambridge University Press.
- [22] M. Kunjir, B. Fain, K., and S. Babu. 2017. ROBUS: Fair Cache Allocation for Data-parallel Workloads. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD)*. 219–234.
- [23] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*. 125–131.
- [24] N. Luzum. 2016. Voters to Decide Fate of Durham Bond Referenda. *The Chronicle* (2016). <http://www.dukechronicle.com/article/2016/10/voters-to-decide-fate-of-durham-bond-referenda>
- [25] T. 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.
- [26] J. F. Nash. 1950. The Bargaining Problem. *Econometrica* 18, 2 (1950), pp. 155–162.
- [27] D. C. Parkes, A. D. Procaccia, and N. Shah. 2012. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC ’12)*. 808–825.
- [28] PBP. 2015. Participatory Budgeting Project. Website. <http://www.participatorybudgeting.org/>
- [29] B. Plaut and T. Roughgarden. 2018. Almost Envy-freeness with General Valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2584–2603.
- [30] S. A. Plotkin, D. B. Shmoys, and É. Tardos. 1995. Fast Approximation Algorithms for Fractional Packing and Covering Problems. *Math. Oper. Res.* 20, 2 (1995), 257–301.
- [31] A. D. Procaccia. 2013. Cake Cutting: Not Just Child’s Play. *Commun. ACM* 56, 7 (July 2013), 78–87.
- [32] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 670–676.
- [33] H. E. Scarf. 1967. The Core of an N Person Game. *Econometrica* 35, 1 (1967), pp. 50–69.
- [34] J. Schummer. 1997. Strategy-proofness Versus Efficiency on Restricted Domains of Exchange Economies. *Social Choice and Welfare* 14 (1997), 47–56.
- [35] T. N. Thiele. 1895. Om flerfoldsvalg. In *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*. 415–441.
- [36] H. R. Varian. 1976. Two problems in the theory of fairness. *Journal of Public Economics* 5, 3-4 (1976), 249–260.